

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE 31 January 1997	3. REPORT TYPE AND DATES COVERED Final Technical Report 01 Sept 92-30 Nov 95
4. TITLE AND SUBTITLE Wavelet-Unifed Design Tool System for Digital Signal Processing			5. FUNDING NUMBERS F49620-92-C-0054
6. AUTHOR(S) Dr. Peter Niels Heller and Dr. Howard L. Resnikoff			
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Aware, Inc. One Oak Park Bedford, MA 01730 AFOSR-TR-97 0313			
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) USAF, AFMC Air Force Office of Scientific Research 110 Duncan Ave., Suite B115 Bolling AFB, DC 20332-0001 nm			10. SPONSORING/MONITORING AGENCY REPORT NUMBER
11. SUPPLEMENTARY NOTES			
12a. DISTRIBUTION AVAILABILITY STATEMENT Approved for public release Distribution is unlimited 19971002 038			
13. ABSTRACT (Maximum 200 words) The mathematical theory and engineering applications of rank M wavelets and M -band filter banks are explored. Techniques for wavelet filter design are given, as well as the theory of smoothness (regularity) for wavelet systems. An overview of the WaveTool signal processing software package follows. Applications of rank M wavelets to telecommunications and data compression are described. Theoretical and laboratory results concerning wavelet-based multicarrier modulation for last-mile communications are reviewed. In addition, improvements to wavelet image compression and new applications to seismic, multispectral, and sonar data are presented. Finally, signal processing with multiwavelets is introduced. The interaction between the theory of rank M wavelets and the requirements of real-world applications is considered throughout. DTIC QUALITY INSPECTED 4			
14. SUBJECT TERMS wavelets, signal processing, digital modulation, data compression			15. NUMBER OF PAGES 100
			16. PRICE CODE
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT UL

Wavelet-Unified Design Toolbox for Digital Signal
Processing
Final Technical Report
ARPA/AFOSR Contract F49620-92-C-0054

Principal Investigators:
Peter Niels Heller and Howard L. Resnikoff
Aware, Inc.
One Oak Park
Bedford, MA 01730

January 31, 1997

Distribution Statement A. Approved for public release; distribution is unlimited.

Effort sponsored by the Air Force Office of Scientific Research, Air Force Materiel Command, USAF, under contract number F49620-92-C-0054. The U. S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon.

The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the Air Force Office of Scientific Research or the U. S. Government.

Contents

1	Executive Summary	3
2	Project Goals	4
3	Review of the Rank M Wavelet Framework	5
3.1	Rank M Wavelet Matrices and Filter Banks	5
3.2	Rank M Wavelets and Multiscale Decompositions	10
4	Theoretical Results	12
4.1	Structure of Rank M Wavelet Systems	13
4.2	Vanishing Moments and Polynomial Approximation for Rank M Wavelets	15
4.3	Smoothness (Regularity) of Rank M Wavelets	16
4.4	Methods for Wavelet Filter Design	23
4.5	Construction of Full Rank M Wavelet Matrices	42
5	WaveTool Software	50
6	Applications	55
6.1	DWMT – Multicarrier Modulation via Rank M Wavelets	55
6.2	Image Compression	64
6.3	Sonar Data Compression in Real Time	72
6.4	Multiwavelet Signal Processing	73
7	Technical Conclusions	81
8	Participants	82
9	Publications	82
10	Transition of Technology to Government and Commercial Uses	85
10.1	WaveTool	85
10.2	Other commercial products	85
10.3	Other government work	86

1 Executive Summary

This is the final technical report for the ARPA/AFOSR-sponsored project “A Wavelet-Unified Design Tool System for Signal Processing,” contract F49620-92-C-0054. This project explored the theory of a broad new class of mathematical transforms (rank M wavelets), developed a production-quality software package for the design and implementation of these transforms, and applied them to concrete problems in signal processing and communications. Significant results of the three-year project include:

- Complete theoretical treatment of rank M wavelets, including parametrization, design and implementation techniques. Original work on the smoothness (regularity) of wavelet systems and its application to image compression.
- Development of a broad and flexible software system, “WaveTool”, for designing and implementing wavelet algorithms. After installation at a number of government, academic, and industrial beta sites, this software was successfully turned into a commercial product.
- Application of rank M wavelets and design techniques to broadband communications, particularly in multicarrier modulation. Wavelet transforms developed under this contract have been implemented in a chipset product for high-bitrate last-mile telecommunications.
- Application of the wavelet transforms and design techniques to the compression of sonar, seismic, and multispectral image data, as well as improvements to wavelet-based still image compression.
- First significant applications of emerging multiwavelet techniques to signal and image processing.

The report that follows details both our vision for the project and its substantive achievements, in the hope that others may take the work even further.

2 Project Goals

Wavelet transforms, discovered in the late 1980's, offer a variety of useful properties and important new applications. The current project sprung from the development, at Aware and elsewhere, of a broad new class of transforms called "rank M wavelets." These new methods encompass traditional discrete block transforms such as the Discrete Fourier Transform (DFT), Discrete Cosine Transform and Walsh-Hadamard transforms, as well as Daubechies wavelets and the Lapped Orthogonal Transform (LOT) and Modulated Lapped Transform (MLT) of subband coding. In addition, a wide range of entirely new data transforms are described by the rank M framework, with many potential applications. The goals of the project were threefold:

- To gain a deeper theoretical understanding of the space of rank M wavelets, in both the discrete and continuous domains;
- To build a wavelet software toolbox (known as WaveTool) that enables the engineer to prototype and test the new wavelet methods at his desktop;
- To develop rank M wavelet approaches to new application problems, such as seismic and sonar signal processing, and communications.

We achieved significant results in each of these three areas. Initial work concentrated on answering a number of theoretical questions regarding rank M wavelets, resulting in complete parametrizations, closed-form constructions, and an analysis of the regularity (smoothness) of these function systems. We completed development of the WaveTool wavelet software toolbox, a GUI-based system for UNIX and PC platforms. An engineer can use WaveTool to rapidly design a variety of wavelet filter banks and graphically assemble them into arbitrary tree structures. He or she can then apply the resulting transform to data in a variety of formats, and work with the results both visually and numerically. The WaveTool software package was released to a number of government, academic, and industrial "beta" sites, and, with further modifications, grew into a commercial product offered by Aware. In addition to the effort put into theory and software, we applied rank M wavelet methods to a number of application areas, including the compression of multispectral, seismic, and sonar data. A surprise development was the application of rank M wavelets to multicarrier modulation for broadband last-mile telecommunications. Wavelet filter designs from the WaveTool software proved to be a critical enabler for this "wavelet multitone" technology. Finally, we applied project resources to the development of new "multi-wavelet" methods for signal and image processing (denoising and data compression); this work represents the first significant application of multiwavelets, with promising results.

In the sections to follow we give a background on rank M wavelets and filter banks, followed by sections devoted to the results of the project in the three different directions of theory, software, and applications.

3 Review of the Rank M Wavelet Framework

In this section we review the general framework provided by rank M wavelets, establishing connections with the theory of multirate filter banks and with multiscale analysis and wavelets. This will help set the stage for the discussion of our theoretical and applied work to follow.

Rank M wavelet theory provides a unified methodology for the design and analysis of discrete transforms; its instances share the following important properties:

- The transforms satisfy a perfect-reconstruction (invertibility) property while yielding a critically sampled data representation;
- It is possible to impose orthogonality on the transform basis functions so that the wavelet transform preserves signal energy;
- The transform basis functions have compact support and are represented by digital FIR filters;
- The class includes basis functions and associated filters that correspond to overlapped windows;
- The class includes multirate (polyphase) digital filters;
- The transforms describe a signal in time, frequency, and scale;
- They have “fast” computational algorithms.

This class of data transforms includes and generalizes orthogonal block transforms such as the FFT, DCT, and Hadamard transforms, filter banks such as the Lapped Orthogonal Transform, and the Daubechies wavelets.

3.1 Rank M Wavelet Matrices and Filter Banks

The basic building block of rank M wavelet theory is a wavelet matrix or filter bank, which partitions a signal into M frequency channels (M is an integer greater than or equal to 2). This partitioning, which is achieved by convolution with the matrix rows followed by downsampling, is shown in Figure 1. We insist on orthogonality of the transform, i.e. that the energy in the transform domain is equal to the signal energy, and that we are able to invert the transform using the same matrix. If the $M \times Mg$

wavelet matrix \mathbf{A} has entries $a_{r,k}$, so that each row of the matrix is a filter impulse response, then the orthogonality condition can be stated as

$$\sum_k a_{r,k} a_{r',k+Ml} = M \delta_{r,r'} \delta_{l,0} . \quad (1)$$

In other words, the matrix of filter taps is orthogonal to itself when overlapped at shifts which are multiples of M . A key consideration here is that the *genus* g of the matrix may be greater than 1, that is, the transform may have nontrivial overlap (when $g = 1$, a wavelet matrix is simply an $M \times M$ orthogonal transform, such as the DFT). Each input “frame” of M samples contributes to more than one output frame, when $g > 1$. This extension from block transforms to overlapping transforms is important, for it enables the signal processor to optimize the filters in the transform for additional properties such as high stopband attenuation (better subchannelization), or regularity (smoothness) in a multiresolution tree structure.

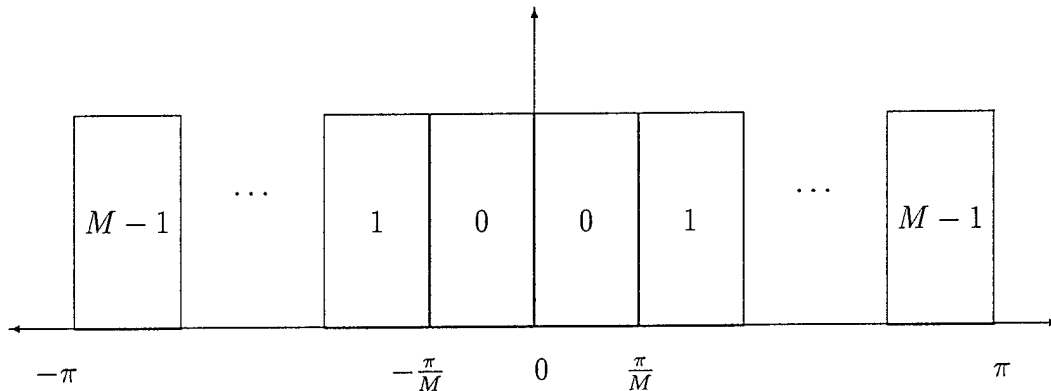


Figure 1: Ideal frequency responses of an M -band filter bank (wavelet matrix).

Table 1 indicates where various conventional digital signal processing transforms fit into the general rank M framework, as the wavelet matrix parameters M and g vary. Every box in the diagram describes a family of filters, some of which generalize classical filters, others of which are entirely new. The top row includes the rank 2 Daubechies wavelets of compact support, while the first column encompasses classical block transforms. The interior of the table, where $M > 2$ and $g > 1$, includes early examples like the Lapped Orthogonal Transform and the Modulated Lapped Transform, but also much uncharted territory. One of the goals of this project has been to provide design tools for creating wavelet transforms of arbitrary rank M and genus g , and pushing them into new applications.

Many transform application domains fit within the context of the arbitrary rank wavelet theory. For example, rank 2 overlap 3 wavelets have been found to be optimal for image and video compression [130], while the JPEG image compression algorithm

Table 1: Conventional digital signal processing transforms within the framework of the arbitrary rank wavelet theory. $\mathcal{W}(M, g)$ denotes the set of wavelet transforms of rank M and genus (overlap) g , while Dg =Daubechies' wavelet transform with $2g$ coefficients.

	$g = 1$	2	3	...	g	...
$M = 2$	Haar	$D2$	$D3$...	Dg	...
\vdots	\vdots				\vdots	
8	8-pt DCT	rank 8 LOT and MLT			$\mathcal{W}(8, g)$	
	8-pt FFT					
	8-pt Hadamard					
	8-pt Chebyshev					
	8-pt Slant					
	\vdots					
	$\mathcal{W}(8, 1)$					
\vdots	\vdots				\vdots	
M	$\mathcal{W}(M, 1)$				$\mathcal{W}(M, g)$	
\vdots	\vdots				\vdots	

is based on a rank 8 overlap 1 DCT block transform [75]. The MPEG Layer II audio compression standard [5] employs a rank 32 genus 16 wavelet transform. In error control coding [13], Reed-Solomon codes are based on the overlap 1 (finite field) FFT and Hadamard block codes are based on the overlap 1 Hadamard matrices. One of the greatest successes of this project has been the application of wavelet matrix methods to multicarrier modulation for broadband communications. While the ANSI T1E1.4 multicarrier standard for Asymmetric Digital Subscriber Line (ADSL) transmission [6] specifies the use of a rank 512 genus 1 Discrete Fourier Transform, Aware's submissions for the next-generation Very high bitrate Digital Subscriber Line (VDSL) standard are based on a rank 128 genus 6 wavelet transform [109], [89], [47].

Every rank M wavelet matrix can be realized as an M -band multirate filter bank [116], with each row of the matrix corresponding to a filter in the filter bank. A multirate filter bank works as shown in Figure 2: a signal is passed through a bank of M different filters (and these filters usually partition the frequency spectrum), and the M results are downsampled (or decimated) by a factor of M , so that while each output v_k (called a subband) has a sampling rate $\frac{1}{M}$ that of the original signal, the totality of the outputs retains the critical sampling rate. This filtering-and-downsampling operation is the fundamental building block of a rank M wavelet transform. The transform or subband domain provides a new representation of the signal, in which

processing such as quantization (for compression) or thresholding (for denoising) or peak-picking (for feature extraction) may take place. In order to invert the transform, the subbands are upsampled (interpolated) by a factor of M , passed through the (time-reversed) filters of the wavelet matrix, and summed to yield the reconstructed signal. If no processing was done in the transform domain, condition (1) ensures that the reconstructed signal is identical to the input, up to a delay.

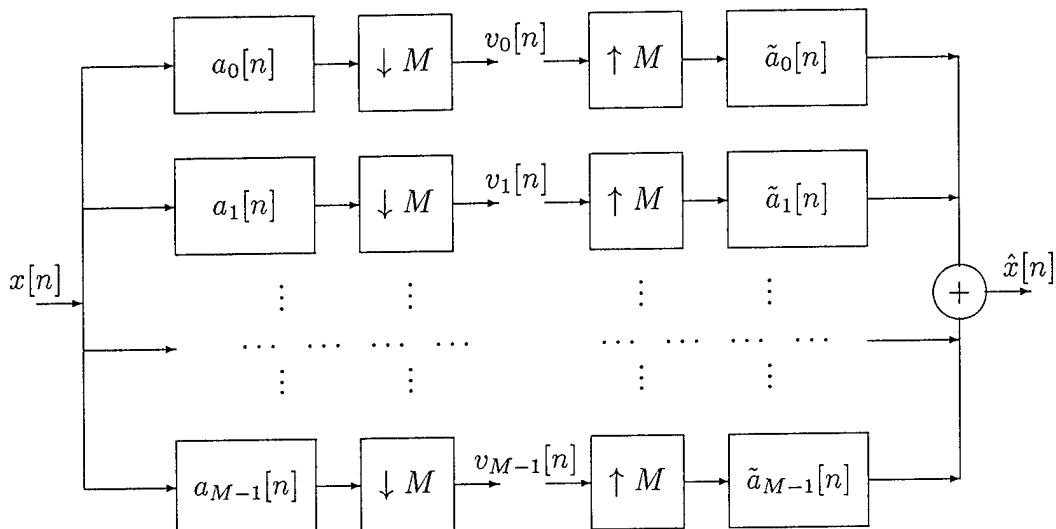


Figure 2: An M -band multirate filter bank associated with a wavelet matrix

We consider briefly one of the benefits of allowing arbitrary overlap or genus in a rank M wavelet transform. The frequency responses $A_k(\omega)$ of the M filters in a wavelet matrix filter bank are designed to uniformly partition the frequency spectrum into M pieces, or subbands. If the filters were ideal “brickwall” filters, this frequency partitioning would look like Figure 1. In practice, each wavelet filter is an FIR filter and so it cannot offer such ideal subchannelization. Consider a widely used wavelet matrix of genus 1, the Discrete Cosine Transform (DCT). The DCT is widely used as a real-valued time-to-frequency transform, with the rank of the transform giving the number of frequency bins. Independent of the rank M , the DCT’s frequency bins are non-ideal to the same extent, with sidelobes that are 13 dB below the main lobe of each component filter. This is shown in Figure 3. By moving to a higher genus or overlapped rank M wavelet transform, it is possible to trade duration in time for better subchannelization, while maintaining orthogonality of the transform. Figure 4 shows the frequency responses of such an overlapped wavelet matrix; the example has rank 16 and genus 4, and was constructed with Aware’s WaveTool software. Each tone now has sidelobes that lie 35 dB below the main lobe, providing a superior approximation to the ideal filter bank of Figure 1.

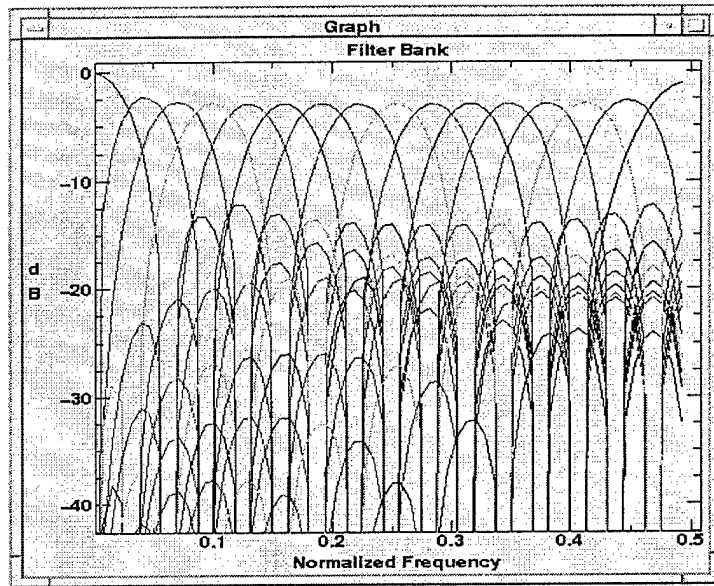


Figure 3: Frequency responses of the filters in a rank 16 Discrete Cosine Transform.

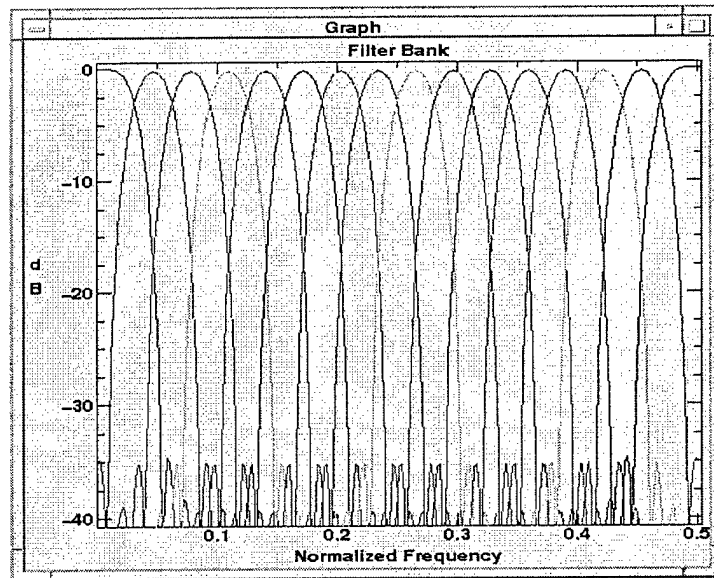


Figure 4: Frequency responses of the filters in a rank 16 genus 4 wavelet matrix constructed using the WaveTool software.

As we will see below, the possibility of using overlapped transforms also makes possible the construction of *wavelets* (multiresolution analyses) with both compact support and underlying smoothness. This was originally seen by Daubechies [23], and the extension to arbitrary values of M was one of the goals and accomplishments of this research project.

3.2 Rank M Wavelets and Multiscale Decompositions

The essence of the wavelet idea is to assemble such building blocks into a tree structure, providing a decomposition of a signal into multiple resolutions, with nonuniform frequency partitioning. The original such wavelet decomposition was the Mallat tree [64], which iterates on the lowpass output of a rank two wavelet building block to create an octave-band decomposition (Figure 5). Such an octave-band decomposition matches the human visual system, and so is well-suited for transform coding of images.

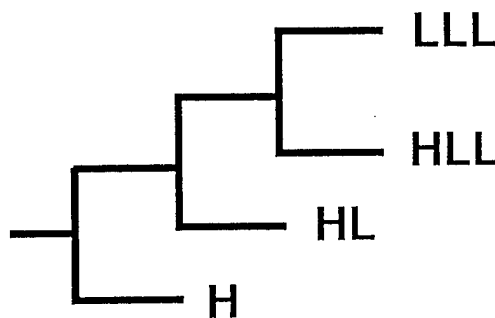


Figure 5: 3-level Mallat tree made up of rank 2 wavelets

For this rank 2 tree, Mallat [64] and Daubechies [23] established a precise mathematical connection between the operation of the two-channel wavelet filter bank and an underlying continuous-time basis of $L^2(\mathbf{R})$. In particular, they found that the imposition of additional linear conditions on the filter bank or wavelet matrix enabled the Mallat iteration to converge to a basis function for continuous-time. Not only is this condition necessary for the infinite iteration to make sense; it even makes a difference for a wavelet tree decomposition of a finite size [87]. Figure 6 compares the overall impulse response that results from iterating a Smith-Barnwell 8-tap orthogonal non-wavelet filter eight times on the lowpass output with the impulse response of the same iteration using a Daubechies 8-tap wavelet filter. The iterated Smith-Barnwell filter displays a fractal behavior which makes it inappropriate for use in a tree-structured decomposition. Further “regularity” or “vanishing moment” constraints on the low-pass filters lead to smoother continuous-time bases, and smoother finite iterates (the

Daubechies 8-tap filter used in Figure 6 has regularity order 4). Whereas overlapping filter banks (genus $g > 1$) were originally used to obtain better frequency localization than block transforms, in the wavelet setting overlap can be used to obtain smoothness of the iterated filter bank. We have generalized the vanishing moment conditions to the rank M case [100], [39] and developed explicit formulae for the construction of rank M wavelets satisfying N -th order regularity conditions. Furthermore we have discovered a set of results relating these discrete-time conditions to the smoothness (Sobolev exponent) of the corresponding continuous-time basis functions. In particular, we have described the asymptotic regularity of several different constructions for rank M wavelet bases as the filter length (support) increases. These issues are discussed in detail in section 4.3. Given a means for measuring smoothness, it can be turned into a criterion for wavelet filter design, a technique we have exploited in the rank 2 and rank M cases (sections 4.4.3 – 4.4.6).

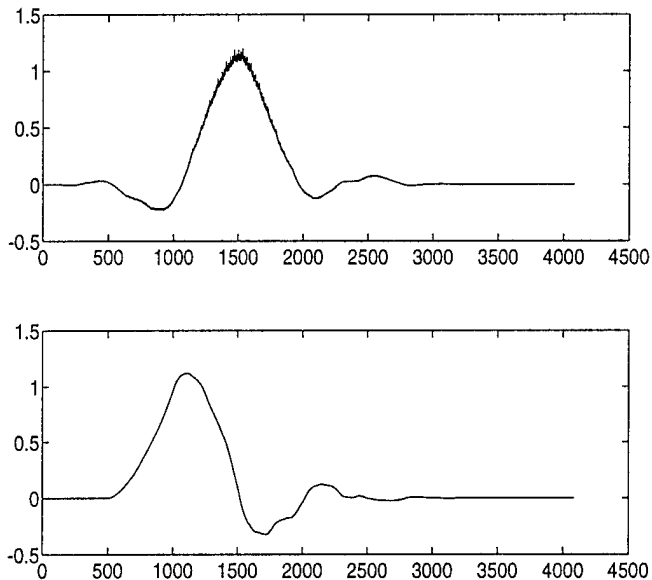


Figure 6: Top: Impulse response of upsampling and convolution with the Smith-Barnwell 8-tap filter, iterated eight times. Bottom: Impulse response of upsampling and convolution with the Daubechies 8-tap wavelet filter, iterated eight times.

The notion of assembling wavelet filter banks into *arbitrary* tree structures, not just Mallat trees which iterate on the lowpass alone, is an essential part of our approach. Arbitrary trees assembled out of rank M wavelets enable a wide range of nonuniform frequency decompositions. Aware's CD-quality real-time audio compression algorithm [98] relies on a two-level tree beginning with a rank 8 filter bank, and iterating on 5 of the 8 outputs with wavelets of varying rank; this tree is depicted in

Figure 7. This higher rank wavelet decomposition produces a frequency/scale partitioning which matches the human ear – a distinctly different decomposition from what one would use for natural images, or seismic data. Each application will have its own best decomposition. Understanding and applying rank M multiresolution decompositions, as well as building a set of software tools to help propagate these techniques, have been the goals for the Higher Rank Wavelet project.

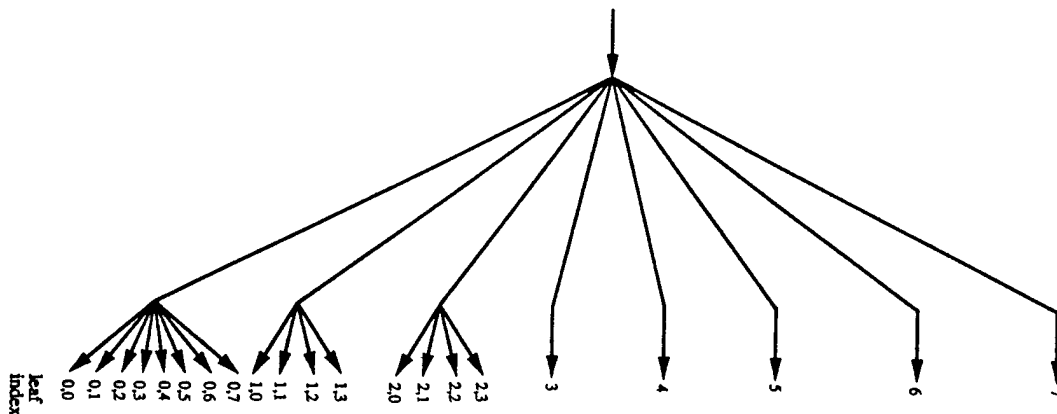


Figure 7: Tree used for decomposition of audio signals in Aware’s real-time CD-quality compression algorithm.

4 Theoretical Results

In this first of three sections reviewing the technical results of the Higher Rank Wavelet Toolbox project, we examine the theoretical results on higher rank wavelets that were obtained. The structure and parametrization of rank M discrete wavelets are discussed, as well as means for measuring the regularity of the associated continuous-time bases. We then cover a number of different techniques that have been developed for rank M wavelet filter design, or in mathematical terms, the construction and optimization of finite rank M wavelet sequences with desirable properties such as subchannelization or smoothness of the associated continuous-time wavelet basis.

4.1 Structure of Rank M Wavelet Systems

A discrete rank M wavelet system [35, 46, 125, 131] has one scaling sequence (lowpass filter) and $M - 1$ wavelet sequences (bandpass and highpass filters). A real rank M wavelet system is given by an $M \times K$ *wavelet matrix* \mathbf{A} whose entries $a_{s,k}$ satisfy

$$\sum_k a_{s,k} a_{s',k+Ml} = M \delta_{s,s'} \delta_{0,l} , \quad (2)$$

$$\text{and } \sum_k a_{s,k} = M \delta_{s,0} . \quad (3)$$

The rows of \mathbf{A} are orthogonal under shifts of M . In contrast to the rank 2 case, where the single wavelet sequence is determined by the scaling sequence, the rank M case has considerable freedom in the choice of the $M - 1$ wavelet sequences. The construction of rank M wavelets can be broken into two steps: the design of a lowpass filter (scaling sequence $a_{0,k}$) and the construction of a full rank M wavelet matrix, given the scaling sequence and other information. We describe several methods for the design of wavelet lowpass filters in Section 4.4, and approaches to wavelet matrix construction in Section 4.5.

4.1.1 Wavelet Bases

As observed in [46], [96], the wavelet matrix describes a basis for the space of functions on \mathbf{Z} . Specifically, a discrete function $f(k)$ may be expanded

$$f(k) = \sum_{s=0}^{M-1} \sum_{l=-\infty}^{\infty} c_{s,l} a_{s,Ml+k} \quad (4)$$

where

$$c_{s,l} = \frac{1}{M} \sum_k f(k) a_{s,Ml+k} . \quad (5)$$

This discrete basis property is equivalent to the orthogonality condition (2) – the wavelet matrix provides a set of *overlapping basis functions* for $\ell^2(\mathbf{Z})$. The additional “low-pass” condition (3) for wavelets confers the ability to develop orthonormal bases of $L^2(\mathbf{R})$.

It is natural to take these rank M discrete wavelet systems and seek to construct compactly supported rank M wavelet orthonormal bases of $L^2(\mathbf{R})$; initial steps in this direction have been taken by Gopinath and Burrus [35]. One begins by considering the rank M scaling function ϕ associated with the scaling sequence $a_{0,k}$, which is the solution to the dilation equation

$$\phi(x) = \sum_k a_{0,k} \phi(Mx - k) . \quad (6)$$

Thus there is a one-one correspondence between the scaling sequences discussed here and scaling functions on \mathbf{R} . The dilation equation (6) has a unique compactly supported solution in $L^2(\mathbf{R}) \cap L^1(\mathbf{R})$; its Fourier transform $\hat{\phi}$ satisfies

$$\hat{\phi}(\xi) = \prod_{j=1}^{\infty} A_0(\xi/M^j) .$$

The corresponding family of wavelets $\psi_{r,j,k}$ ($1 \leq r \leq M-1$) is defined from ϕ by

$$\psi_r(x) = \sum_k a_{r,k} \phi(Mx - k)$$

and

$$\psi_{r,j,k} = M^{-j/2} \psi_r(M^{-j}x - k) .$$

Lawton [60] proved that the collection $\{\psi_{r,j,k}(x)\}$ form a tight frame for $L^2(\mathbf{R})$.

The explicit formulae for scaling sequences developed below enable one to construct rank M wavelet bases with arbitrary smoothness (as measured by Sobolev differentiability), as discussed in section 4.3.

4.1.2 The Polyphase Matrix Representation

Matrices satisfying the orthogonality condition (2) have been extensively studied in the signal processing literature [113, 116, 118] under the name “ M -band paraunitary perfect reconstruction filter banks.” Engineers often restate (2) in the z -transform domain, as follows: form the $M \times M$ *polyphase matrix* $\mathbf{H}(z)$ with polynomial entries

$$h_{s,r}(z) = \sum_l a_{s,r+lM} z^l .$$

Observe that \mathbf{H} and \mathbf{A} are related by

$$\mathbf{H}(z) = \mathbf{A}_0 + z\mathbf{A}_1 + \dots + z^{g-1}\mathbf{A}_{g-1} .$$

$\mathbf{H}(z)$ is said to be *paraunitary* if $\frac{1}{\sqrt{M}}\mathbf{H}(z)$ is unitary on the unit circle:

$$\mathbf{H}(z)\mathbf{H}^\dagger(z^{-1}) = M\mathbf{I} , \text{ for } |z| = 1 . \quad (7)$$

Comparison of coefficients of powers of z shows that the paraunitarity of \mathbf{H} is equivalent to the orthogonality under shifts (2) of the wavelet system. Paraunitary matrices and polyphase factorizations have been investigated in great detail [116]. We impose the additional linear condition (3) to form a wavelet matrix – this amounts to the requirement that the matrix $\mathbf{H}(z)|_{z=1} = \mathbf{H}_0$ be a Haar wavelet matrix. This proves essential for the development of orthonormal bases of $L^2(\mathbf{R})$ (cf. [23, 35]).

The orthogonality condition (2) can also be stated in the Fourier domain – for each of the sequences a_s , consider its “symbol” or Fourier transform

$$A_s(e^{i\omega}) = \frac{1}{M} \sum_{k=0}^{K-1} a_{s,k} e^{ik\omega}, \quad 0 \leq s < M.$$

Then (2) is equivalent to

$$\sum_{m=0}^{M-1} A_s(e^{i(\omega+2\pi m/M)}) \overline{A_{s'}(e^{i(\omega+2\pi m/M)})} \equiv \delta_{s,s'}. \quad (8)$$

The paraunitarity (7) of $\mathbf{H}(z)$ implies that

$$\sum_{s=0}^{M-1} |A_s(e^{i\omega})|^2 \equiv 1. \quad (9)$$

4.2 Vanishing Moments and Polynomial Approximation for Rank M Wavelets

The notion of *vanishing moments* of a wavelet systems was originally introduced by Daubechies in the rank 2 case [23] as a means of ensuring differentiability of the continuous-time basis functions. Since then, vanishing moments have come to be seen as important for discrete-time decompositions as well, and to offer important information on the power with which wavelet representations can approximate polynomials [101]. As in the rank 2 case, there are several equivalent definitions for a rank M wavelet matrix $\mathbf{A} = \{a_{s,k}\}$ to have N vanishing moments, which we have identified [100, 39] in the following theorem:

Theorem 4.1 *A rank M wavelet system is said to have approximation degree N if any of the following equivalent conditions hold:*

- (i) $(d/d\xi)^n \hat{\psi}_r(0) = 0$, for $n = 0, 1, \dots, N-1$ and $r = 1, 2, \dots, M-1$.
- (ii) $(d/d\omega)^n A_r(0) = 0$, for $n = 0, 1, \dots, N-1$ and $r = 1, 2, \dots, M-1$.
- (iii) $(d/d\omega)^n A_0\left(\frac{2\pi k}{M}\right) = 0$, for $n = 0, 1, \dots, N-1$; $k = 1, 2, \dots, M-1$.
- (iv) $\sum_l (k + Ml)^n a_{0,k+Ml}$ is independent of k for $k = 0, 1, \dots, M-1$ and $n = 0, 1, \dots, N-1$.
- (v) $A_0(\omega) = (1 + e^{i\omega} + \dots + e^{i(M-1)\omega})^N Q(\omega)$ for some trigonometric polynomial $Q(\omega)$.

Strang [101] gives additional formulations of “degree N ,” relating it to approximation of order N for functions on $L^2(\mathbf{R})$, i.e. in the continuous-time domain.

These vanishing moment conditions may be used in the design of wavelet lowpass filters; indeed, the definition and systematic exploitation of rank M vanishing moment constraints is described in Section 4.4 below.

4.3 Smoothness (Regularity) of Rank M Wavelets

Once one has constructed a compactly supported continuous-time wavelet basis from a rank M wavelet matrix, it is natural to examine the regularity (differentiability) of the basis elements $\psi_r(x)$. We have developed mathematical tools for measuring regularity in the rank M setting, and explicit constructions of several families of arbitrarily differentiable rank M wavelet bases. We describe our theoretical results here; the constructions are described in section 4.4 below.

An early theorem on wavelets ([24], p.153) states that if a wavelet $\psi_r(x)$ is N times continuously differentiable and has sufficient decay to be in $L^1(\mathbf{R})$, then $\widehat{\psi_r}$ must vanish to order N at $\xi = 0$. This led to the imposition of vanishing wavelet moments (one of the equivalent conditions of Theorem 4.1) for the construction of wavelet systems. Since the wavelets $\psi_{r,j,k}$ are linear combinations of dilates of the scaling function ϕ , the question of their differentiability reduces to “How differentiable can we make $\phi(x)$?” One standard way to measure differentiability of a function in $L^2(\mathbf{R})$ is via Sobolev spaces, which measure the number of L^2 derivatives a function has.

Given a real number s , the Sobolev space \mathcal{H}^s is defined by

$$\mathcal{H}^s := \left\{ f : \|f\|_{\mathcal{H}^s}^2 = \int_{\mathbf{R}} |\hat{f}(\xi)|^2 (1 + |\xi|^2)^s d\xi < \infty \right\} .$$

Recall that for $f : \mathbf{R} \rightarrow \mathbf{C}$, $f \in \mathcal{H}^s \Rightarrow f \in \mathcal{C}^\alpha$, for $\alpha < s - \frac{1}{2}$. That is, if $s > n + \frac{1}{2}$, then f has n continuous derivatives. We define the *Sobolev smoothness* of a function $\phi \in L^2(\mathbf{R})$ to be

$$s(\phi) := \sup \{s : \phi \in \mathcal{H}^s\} .$$

Daubechies and others [23], [31], [120], [122] developed theoretical tools for measuring the Sobolev exponent of a rank 2 scaling function, and found that the smoothness of a minimal length rank 2 wavelet system with approximation degree N (N vanishing moments) grows linearly with N . Specifically, if ϕ_N denotes the scaling function for such a system,

$$s(\phi_N) \approx .2075N .$$

We set out to generalize this result to the rank M case, and have found surprising results [52].

As in the rank 2 case, the Sobolev regularity of the scaling function ϕ is determined by the maximum eigenvalue of a finite-dimensional operator associated with the scaling sequence $a_{0,k}$. In particular, consider the modulus squared of the lowpass frequency response, denoted by P :

$$P(\omega) = |A_0(\omega)|^2 = \left| \frac{1 + e^{i\omega} + \dots + e^{i(M-1)\omega}}{M} \right|^{2N} R(\omega) . \quad (10)$$

The trigonometric polynomial $R(\omega)$ is defined as the remainder after division of $P(\omega)$ by N powers of the Haar polynomial squared; this remainder will play an important role in both our measure of smoothness and in later filter designs. Once given the functions $P(\omega)$ and $R(\omega)$, we define the associated transfer operators [21] that take functions on $[0, \pi)$ to functions on $[0, \pi)$:

$$\begin{aligned} T_P u(\omega) = & P\left(\frac{\omega}{M}\right) u\left(\frac{\omega}{M}\right) + P\left(\frac{\omega + 2\pi}{M}\right) u\left(\frac{\omega + 2\pi}{M}\right) + \dots \\ & \dots + P\left(\frac{\omega + 2\pi(M-1)}{M}\right) u\left(\frac{\omega + 2\pi(M-1)}{M}\right), \end{aligned}$$

and

$$\begin{aligned} T_R u(\omega) = & R\left(\frac{\omega}{M}\right) u\left(\frac{\omega}{M}\right) + R\left(\frac{\omega + 2\pi}{M}\right) u\left(\frac{\omega + 2\pi}{M}\right) + \dots \\ & \dots + R\left(\frac{\omega + 2\pi(M-1)}{M}\right) u\left(\frac{\omega + 2\pi(M-1)}{M}\right). \end{aligned}$$

Here we have extended the domain of definition of u by its natural periodization.

While they are defined on the infinite-dimensional space $\mathcal{C}[0, \pi]$, these transfer operators leave certain finite-dimensional subspaces invariant, and their finite-dimensional restrictions supply enough information to measure the Sobolev regularity of ϕ . We now sketch how this is done. Let the trigonometric polynomials P and R be given as in (10), expressed in the form

$$P(e^{i\omega}) = \sum_{j=-L}^L p_j e^{ij\omega}, \quad R(e^{i\omega}) = \sum_{j=-J}^J \rho_j e^{ij\omega}, \quad (11)$$

Moreover define the $2L + 1$ -dimensional space of trigonometric polynomials

$$\mathcal{E}_L := \left\{ u(\omega) = \sum_{j=-L}^L u_j e^{ij\omega} \right\},$$

the analogous $2J + 1$ -dimensional subspace \mathcal{E}_J , and the distinguished subspace

$$\mathcal{F}_{L,N} := \left\{ f \in \mathcal{E}_L : f(\omega) = (1 - \cos \omega)^N g(\omega) \right\}$$

of \mathcal{E}_L . We have demonstrated the following theorem, generalizing the rank 2 work of [24], [31], [120]:

Theorem 4.2 *Suppose that the wavelet system associated with the scaling sequence $\{a_k\}$ has degree N (the first N wavelet moments vanish). Then*

(i) The subspace \mathcal{E}_L of $\mathcal{C}[0, \pi]$ is an invariant subspace for the operator T_P , i.e.

$$T_P : \mathcal{E}_L \rightarrow \mathcal{E}_L .$$

The operator T_P acting on \mathcal{E}_L has the matrix representation $[T_P]_{j,k} = p_{Mj-k}$, with respect to the basis $\{e^{-iL\omega}, \dots, 1, \dots, e^{iL\omega}\}$, where $p_l = 0$ if $l < -L$ or $l > L$. This matrix is pseudo-circulant, obtained from shifts of the sequence p_l by M for each row. For example, if $L = Mg - 1$ for some $g \in \mathbb{Z}$, as is the case when $\{a_k\}$ has length Mg , the matrix has the form

$$T_P = \begin{bmatrix} 0 & 0 & & \dots & & 0 & 0 \\ & & & \ddots & & & \\ 0 & 0 & & \dots & & 0 & 0 \\ p_{L-M+2} & \dots & p_L & 0 & 0 & \dots & 0 & 0 \\ p_{L-2M+2} & \dots & p_{L-M} & p_{L-M+1} & \dots & p_L & 0 & \dots & 0 \\ & & & \ddots & & & & & \\ p_{-L} & \dots & p_{M-2-L} & p_{M-1-L} & p_{M-L} & \dots & & p_L \\ 0 & \dots & 0 & 0 & p_{-L} & \dots & & p_{L-M} \\ & & & \ddots & & & & \\ 0 & 0 & & \dots & 0 & p_{-L} & \dots & p_{M-2-L} \\ 0 & 0 & & \dots & & & 0 & 0 \\ & & & & \ddots & & & \\ 0 & 0 & & \dots & & & 0 & 0 \end{bmatrix} .$$

(ii) Similarly, the subspace \mathcal{E}_J of $\mathcal{C}[0, \pi]$ is an invariant subspace for the operator T_R , i.e.

$$T_R : \mathcal{E}_J \rightarrow \mathcal{E}_J .$$

We write \tilde{T}_R for the restriction of T_R to the subspace \mathcal{E}_J ; it has the matrix representation $[\tilde{T}_R]_{j,k} = \rho_{Mj-k}$, with respect to the basis $\{e^{-iJ\omega}, \dots, 1, \dots, e^{iJ\omega}\}$.

Here $\rho_j = 0$ if $j < -J$ or $j > J$. This matrix has a pseudocirculant form similar to that of T_P above.

(iii) $T_P|_{\mathcal{E}_L}$ includes among its eigenvalues the numbers M^{-n} , $n = 0, 1, \dots, 2N - 1$.

(iv) Suppose $f \in \mathcal{F}_{L,N}$, $f(\omega) = (1 - \cos \omega)^N g(\omega)$; then

$$T_P f(\omega) = \frac{1}{M^{2N}} (1 - \cos \omega)^N T_R g(\omega) . \quad (12)$$

Hence $\mathcal{F}_{L,N}$ is an invariant subspace for T_P , and moreover if f is an eigenvector for T_P in $\mathcal{F}_{L,N}$ with eigenvalue λ , then

$$T_R g(\omega) = M^{2N} \lambda g(\omega),$$

so that g is an eigenfunction for T_R with eigenvalue $M^{2N} \lambda$.

(v) If $\mu(T_P)$ is the largest eigenvalue of $T_P|_{\mathcal{F}_{L,N}}$ and $\mu(T_P) < \frac{1}{2}$, then the Sobolev exponent of ϕ is

$$s(\phi) = -\frac{\log \mu(T_P)}{2 \log M}. \quad (13)$$

(vi) If $\mu(T_R)$ is the largest eigenvalue of $T_R|_{\mathcal{F}_{L,N}}$ and $\mu(T_R) < \frac{M^{2N}}{2}$,

$$s(\phi) = N - \frac{\log \mu(T_R)}{2 \log M}. \quad (14)$$

Thus we have obtained an explicit formula for the Sobolev smoothness of the scaling function ϕ in terms of the maximal eigenvalue of the associated multiresolution operator T_P acting on a distinguished subspace of the finite-dimensional space \mathcal{E}_{N-1} , and a corresponding formula in terms of the maximal eigenvalue of T_R . These formulae can be used both to measure the Sobolev smoothness of a scaling function, given the scaling sequence that defines it, and as a tool for use in filter design. We explore the first of these uses immediately, and leave the second to section 4.4.

Not only is the smoothness determined by the eigenvalue $\mu(T_R)$ or $\mu(T_P)$, it has been shown [31], [52] that one may estimate $\mu(T_R)$ just from the values of the frequency response $R(\omega)$:

Theorem 4.3 *We have*

$$\mu(T_R) \approx R(\omega_c),$$

where

$$\omega_c = \frac{M\pi}{M+1} \text{ for even } M,$$

and

$$\omega_c = \pi \text{ for odd } M,$$

with the approximation improving as N grows large.

This provides a technique for investigating the asymptotic differentiability of rank M minimal length degree N wavelet bases as the degree N approaches infinity. The difference between the even M and odd M cases is critical; we have proven [52]:

Theorem 4.4 *Consider the minimal length trigonometric polynomials yielding N -th order approximation, with remainders R_N defined in (16), (17) and (18) below. For $M = 2$ and 4, the asymptotic value of $R_N(\omega_c)$ as N grows large is*

$$R_N\left(\frac{M\pi}{M+1}\right) \sim M^{2N} \left(\frac{1 - \cos \frac{M\pi}{M+1}}{2}\right)^N \frac{1}{\sqrt{N}}.$$

Thus, if ϕ_N is the associated scaling function,

$$s(\phi_N) \sim N \frac{\log \frac{2}{1 - \cos \frac{M\pi}{M+1}}}{2 \log M}.$$

The growth in smoothness is approximately linear with N , with a slope of $1 - \frac{\log 3}{\log 4} \sim .2075$ when $M = 2$ and $\frac{3}{4} - \frac{\log 5 + \sqrt{5}}{4 \log 2} \sim 0.0362$ when $M = 4$.

For $M = 3$, the asymptotic behavior of $R_N(\omega_c)$ is given by

$$R_N(\pi) \sim M^{2N} \frac{1}{\sqrt{N}},$$

and thus

$$s(\phi_N) \sim \frac{\log N}{4 \log M}.$$

The growth in smoothness with increasing filter length is only logarithmic for $M = 3$.

The logarithmic growth when $M = 3$ is due to the fact that we evaluate $R(\omega)$ at the critical point π rather than at $\frac{M\pi}{M+1}$. We conjecture that the Sobolev exponent of ϕ_N grows as

$$s(\phi_N) \approx \frac{\log N}{4 \log M}$$

for all odd M . In the case of even M , we conjecture that the Sobolev exponent of ϕ_N grows linearly, with slope

$$s_M = \frac{\log \frac{2}{1 + \cos \frac{\pi}{M+1}}}{2 \log M} \sim \frac{\pi^2}{8(M+1)^2 \log M}.$$

This approaches zero quite rapidly as M increases. Numerical measurements of the Sobolev exponents confirm these conjectures for sequence lengths up to 100, and dilation factors $M \leq 6$. One implication of this result is that adding vanishing moments is not a very effective means of increasing the smoothness of rank M wavelet bases for odd M ; improved approaches are explored in sections 4.4.1 and 4.4.6.

The asymptotic growth described in the theorem is shown in Figure 8, which compares the true Sobolev exponents of the minimal length scaling functions (as computed in [52]) with their predicted linear and logarithmic asymptotes. The difference between the linear growth of the cases $M = 2$, $M = 4$ and the case $M = 3$ is highlighted in Figure 9 which compares the graphs of scaling functions for these three dilation factors, with $N = 2$ and $N = 12$.

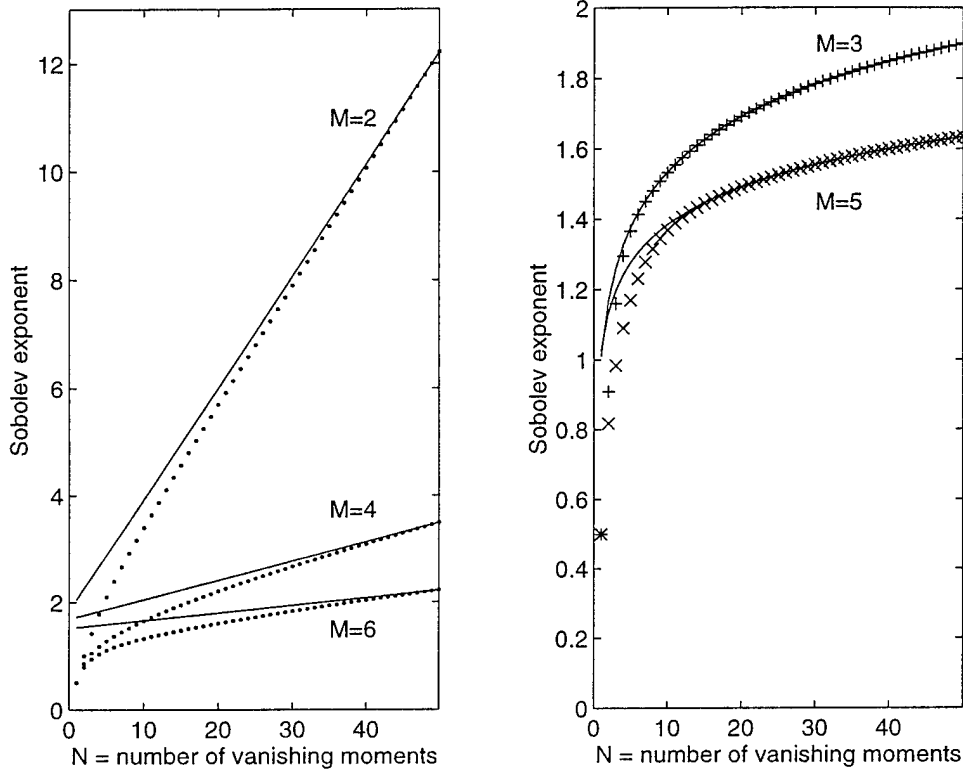


Figure 8: Asymptotic growth in Sobolev smoothness for minimal length degree N scaling functions, $N = 1, 2, \dots, 50$. The left graph shows the actual Sobolev exponents and their linear asymptotes for even rank ($M = 2$, $M = 4$, $M = 6$), while the right graph shows the Sobolev exponents and their logarithmic asymptotes for odd rank ($M = 3$, $M = 5$).

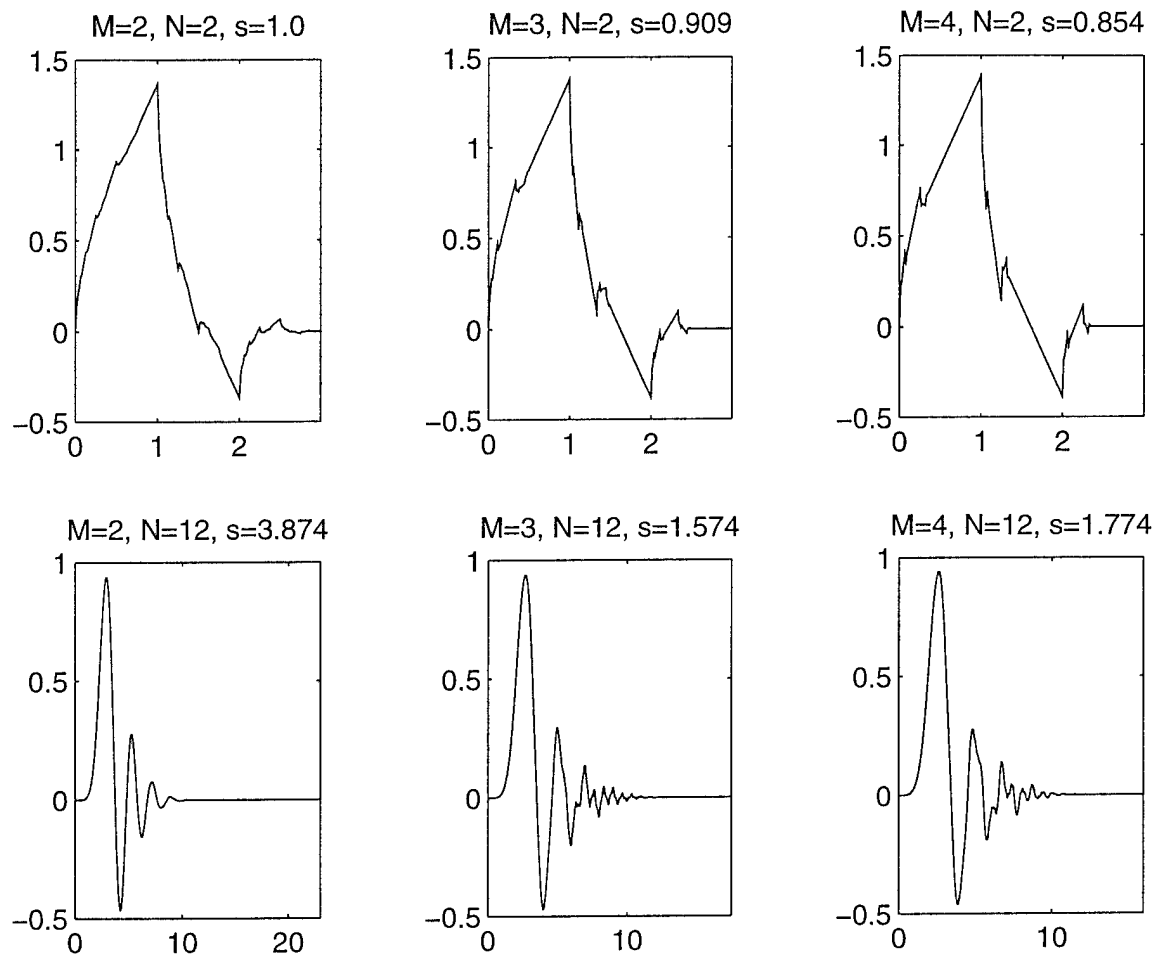


Figure 9: Scaling functions for $M = 2, 3, 4$ and $N = 2$ (top row) and $M = 2, 3, 4$ and $N = 12$ (bottom row).

4.4 Methods for Wavelet Filter Design

In this section we describe a number of methods that we have developed and applied for the design of wavelet lowpass filters (scaling sequences), including the imposition of vanishing wavelet moments (polynomial approximation), smoothness of the wavelet scaling function (both maximally smooth and frequency-sampling-based “higher-smoothness” lowpass filters), and Nguyen’s quadratic-constrained design approach. Additional results appear in [84]

4.4.1 Approximation-based Criteria (Vanishing Moments)

Our first constructions of rank M wavelet lowpass filters were a class that generalize Daubechies orthogonal constructions [23] from the rank 2 case. As in [23], these constructions take off from the characterization of vanishing moments given in Theorem 4.1. Recall from that theorem that the frequency response of an orthogonal wavelet lowpass filter with approximation of degree N (N vanishing wavelet moments) must have the form

$$A_0(\omega) = (1 + e^{i\omega} + \dots + e^{i(M-1)\omega})^N Q(\omega) .$$

Note that $1 + e^{i\omega} + \dots + e^{i(M-1)\omega}$ is the Fourier transform of the Haar scaling sequence $\{1, 1, \dots, 1\}$ of length M . In order to characterize minimal length rank M wavelet systems with N vanishing moments, we shall describe the autocorrelation $|A_0(\omega)|^2$ and then perform a spectral factorization to obtain A_0 . As in the previous section, for a wavelet system with N vanishing moments, we write

$$P(\omega) = |A_0(\omega)|^2 = \left| \frac{1 + e^{i\omega} + \dots + e^{i(M-1)\omega}}{M} \right|^{2N} R(\omega) \quad (15)$$

where we have factored the square modulus $P(\omega)$ into two parts, N powers of the modulus squared Haar polynomial and a term $R(\omega) = |Q(\omega)|^2$ designed to restore orthogonality to the overall symbol.

We have obtained [39] a complete closed-form expression for the minimal length “remainder” term $R_N(\omega)$ for arbitrary values of M and N :

$$R_N(\omega) = \sum_{n=0}^{N-1} r_n (1 - \cos \omega)^n , \quad (16)$$

where

$$r_n = \sum_{k_1+k_2+\dots+k_{M_1}=n} \left\{ \prod_{m=1}^{M_1-1} \binom{2N+k_m-1}{2N-1} \left(1 - \cos \frac{2\pi m}{M} \right)^{-k_m} \right\} \times \binom{N+k_{M_1}-1}{N-1} (1 - \cos \pi)^{-k_{M_1}} \quad (17)$$

for M even, with $M_1 = \frac{M}{2}$. For M odd,

$$r_n = \sum_{k_1+k_2+\dots+k_{M_1}=n} \left\{ \prod_{m=1}^{M_1} \binom{2N+k_m-1}{2N-1} \left(1 - \cos \frac{2\pi m}{M}\right)^{-k_m} \right\}, \quad (18)$$

where $M_1 = \frac{M-1}{2}$. This completely describes the deterministic autocorrelation $P_N(\omega)$ of the minimal length rank M wavelet lowpass filter with regularity order N . For low values of M , the spectral factor $A_0(\omega)$ may be explicitly computed. For example, the $M = 4$, $N = 2$ generalization of Daubechies 4-coefficient scaling sequence is

$$\{a_{0,k}\} = \left\{ \frac{1 \pm \sqrt{11}}{8}, \frac{3 \pm \sqrt{11}}{8}, \frac{5 \pm \sqrt{11}}{8}, \frac{7 \pm \sqrt{11}}{8}, \frac{7 \mp \sqrt{11}}{8}, \frac{5 \mp \sqrt{11}}{8}, \frac{3 \mp \sqrt{11}}{8}, \frac{1 \mp \sqrt{11}}{8} \right\}. \quad (19)$$

The general (non-minimal length) autocorrelation with degree N approximation is of the form

$$P(\omega) = P_N(\omega) + \tilde{P}(\omega) \quad (20)$$

where

$$P_N(\omega) = \left| \frac{1 + e^{i\omega} + \dots + e^{i(M-1)\omega}}{M} \right|^{2N} R_N(\omega)$$

and

$$\tilde{P}(\omega) = |e^{iM\omega} - 1|^{2N} \sum_{n \neq Ml} \tilde{r}_n \cos n\omega,$$

all subject to

$$P(\omega) \geq 0 \text{ for } \omega \in [0, \pi].$$

Any rank M wavelet lowpass filter with N vanishing moments can be obtained as a spectral factor of such a $P(\omega)$, and any such $P(\omega)$ will have finitely many spectral factors. These general parametrizations [39] (see also [24], [126]) can be used to describe various families of orthogonal wavelets, including the smoother families of section 4.4.6, the “coiflets” of [24], and the smoother wavelets of [123].

4.4.2 Lagrange M -th band filters and M -band wavelets

We have also found an independent explicit formula for the minimal length modulus-squared frequency response $P_N(\omega)$, based on M -th band filters and Lagrange interpolation. This approach was introduced in the 2-band case [10, 101] to connect the maximally flat filters of Herrmann [53] with their spectral factors, the Daubechies wavelets [23]. Herrmann’s construction can be generalized to the M -th band case, yielding the minimal length M -th band lowpass filter with $2N$ -th order flatness at $\omega = 0$. The general arbitrary-length solution follows by the methods of the previous

section. Any lowpass filter satisfying the orthogonality condition (2) or (8) will be a spectral factor of an M -th band filter; in fact (8) simply states that $|A_0(\omega)|^2$ is an M -th band filter. In particular, rank M wavelet lowpass filters with N vanishing moments will be spectral factors of the M -th band filter $H(\omega)$ we find below.

Recall that an M -th band filter [67, 115] is a symmetric FIR filter whose impulse response $h[n]$ satisfies

$$h[Mn] = \begin{cases} 1, & n = 0 \\ 0, & \text{otherwise.} \end{cases} \quad (21)$$

When used in an M -fold interpolator, these filters have the useful property of preserving the existing signal samples, while inserting new values in between. Conversely, M -fold decimation of h results in the unit impulse response. When designing FIR M -th band filters, it is desirable to approximate the ideal lowpass filter. When the M -th band filter is used in the construction of an M -band wavelet lowpass filter, it is also desirable to have some regularity – flatness of the frequency response at $\omega = 0$. These are not conflicting requirements, for flatness at $\omega = 0$ together with $H(0) = 1$ forces lowpass characteristics on $H(\omega)$.

The M -th band condition (21) can be restated in the frequency domain:

$$H(\omega) + H\left(\omega + \frac{2\pi}{M}\right) + \dots + H\left(\omega + \frac{2\pi(M-1)}{M}\right) \equiv 1. \quad (22)$$

We find a minimal-length lowpass filter $H(\omega)$ which satisfies this condition, and vanishes to order $2N$ at $\omega = 0$:

$$H(\omega) = 1 + \mathcal{O}(|\omega|^{2N}). \quad (23)$$

We do so by requiring flatness of order $2N$ at the “test frequencies” $2\pi m/M$:

$$H(\omega + 2\pi m/M) = \mathcal{O}(|\omega|^{2N}), \quad 1 \leq m \leq M-1. \quad (24)$$

One solves for these conditions by constructing each of the polyphase components of $H(\omega)$ as a Lagrange interpolate to $e^{ix\omega}$ on the coset

$$\{x_n = Mn + k, n = -N, 1-N, \dots, N-1\}.$$

Our result, proven in detail in [38], is

Theorem 4.5 *The minimal length M -th band lowpass filter with $2N$ -th order flatness is given by*

$$H(\omega) = \frac{1}{M} + \frac{2}{M^2} \sum_{k=1}^{M-1} \sum_{n=0}^{N-1} \prod_{\substack{l=-N \\ l \neq n}}^{N-1} \frac{Ml + k}{l - n} \cos(Mn + k)\omega. \quad (25)$$

In fact, $H(\omega)$ is just the filter $P_N(\omega)$ of the previous section.

Example ($N = 2$) :

$$H(\omega) = \frac{1}{M} + \frac{2}{M^2} \sum_{k=1}^{M-1} (M - k) \cos k\omega . \quad (26)$$

The stopband attenuation of these filters increases with N (for fixed M), as can be seen in the 4-th band examples shown in Figure 10. The M -th band filters given by formula (25) will not approximate the ideal lowpass filter as well as those designed in [67] using numerical optimization. However, the filters given here offer the advantages of explicit formulae and rational coefficients, as well as the flatness at $\omega = 0$ which is essential for the construction of wavelets.

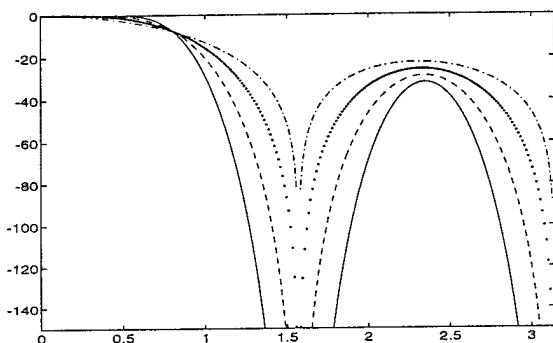


Figure 10: Frequency responses in dB of minimal length 4-th band maxflat filters, $N = 1, 2, 4, 8$.

4.4.3 Maximal Smoothness Filter Design Criteria

The results of section 4.3 lead to a new filter design criterion – we can design FIR lowpass filters which lead to maximally smooth scaling functions (infinitely iterated lowpass filters), as opposed to traditional criteria such as ℓ^2 or ℓ^∞ (Chebyshev) error. Since the smoothness of finite iterates of the lowpass filter qualitatively corresponds to the smoothness of the scaling function [87], we can use the mathematical tools developed for the continuous-time setting to help us design wavelet filters with smooth finite iterates. This is important for applications of wavelets to image coding because the artifacts resulting from lossy quantization reflect the shape of the underlying basis function. If we can optimize the smoothness of the discrete wavelet basis functions resulting from iteration of the lowpass filter, then we should achieve superior perceptual quality for images compressed using these maximally smooth wavelets.

The algorithm for measuring the smoothness of a given lowpass filter is the following:

- Form the deterministic autocorrelation sequence $p[j]$ (time-domain form of $P(\omega)$) and the matrix T_P of Theorem 4.2.
- Find the maximal eigenvalue $\mu(T_P)$ of this matrix, other than the known eigenvalues M^{-n} , $n = 0, 1, \dots, 2N - 1$.
- The Sobolev smoothness (number of L^2 derivatives) of the scaling function associated to the lowpass filter is

$$s(\phi) = -\frac{\log \mu(T_P)}{2 \log M}.$$

An alternative algorithm may be employed, based on the remainder $R(\omega)$ and the reduced transfer operator T_R . Given either method for measuring the wavelet smoothness associated with a lowpass filter, several design approaches may be employed. First, the smoothness measure s may be used as the objective function in an optimization to yield filters that are maximally smooth for a given sequence length; this technique is used in sections 4.4.4 and 4.4.1 below. Secondly, an understanding of the mathematics of Sobolev regularity may be employed to intelligently impose additional zeros on the frequency response of the wavelet lowpass filter, in such a way as to increase the smoothness of the associated scaling function.

4.4.4 Sobolev-optimal Quadrature Mirror Filters

We first applied our Sobolev-optimality criterion to the design of symmetric near-orthogonal Quadrature Mirror Filters (QMF's) [48]. Recall that QMF's are two-channel filter banks (rank 2 wavelets) in which the highpass filter is obtained by an "alternating flip" of the lowpass filter. If the lowpass filter of the pair is $a_0[n]$, a sequence of length N , then the highpass filter is given by

$$a_1[n] = (-1)^n a_0[N - 1 - n]. \quad (27)$$

This can be represented in the frequency domain as

$$A_1(\omega) = e^{i\omega} A_0(\pi - \omega);$$

A_1 is a "mirror" of A_0 . One often combines this construction with orthogonality of the filter bank,

$$\sum_n a_0[n] a_0[n + 2l] = \delta_{0,l} \quad (28)$$

the rank 2 form of equation (2). It was shown rather early in the theory of multirate filter banks [94], [116] that the only symmetric two-channel orthogonal QMF pair is the Haar filter, for which $a_0 = \{\frac{1}{2}, \frac{1}{2}\}$. This filter is too crude for most image coding applications. However, symmetry (linear-phase) for filter banks is an important

property in wavelet and subband image coders; it enables the use of nonexpansive symmetric extension methods at signal boundaries [95]. One solution to this problem is to give up the QMF condition and seek filter banks which are perfect-reconstruction (biorthogonal), rather than paraunitary (orthogonal) [11], [73], [117]. This approach led to the well-known (7,9)-tap biorthogonal wavelet pair used in the FBI's fingerprint compression standard.

A different solution to the problem of symmetry and orthogonality was developed by Adelson and Simoncelli [8]; they *mildly relaxed* the orthogonality condition (28) and by doing so were able to obtain symmetric filters! The odd-length filters they designed satisfied the orthogonality conditions (28) within $1.0\text{e-}3$, and led to near-perfect-reconstruction of wavelet-transformed images. The use of odd-length filters is also noteworthy, for the whole-sample symmetry of an odd-length filter preserves centers of mass, leading to better retention of edge details in a lossy subband compression scheme. Even-length filters will smear single-pixel detail across two filter outputs, leading to smoothed, blurry artifacts.

Several constraints are imposed on the filters: we constrain the length:

$$\sum_n a_n^2 = 1, \quad (29)$$

and we impose one (and therefore two) vanishing wavelet moments:

$$\sum_n a_n = \sqrt{2}, \quad (30)$$

$$\sum_n (-1)^n a_n = 0. \quad (31)$$

Observe that if the filters were truly orthogonal (i.e. fully satisfied the conditions (28)), then one of the equations (30)-(31) would be superfluous. However, in the near-orthogonal case of symmetric QMF's, both equations are necessary. This leaves two free parameters for a nine-tap QMF (one parameter for the seven-tap case). We optimized to find the filter with the greatest Sobolev smoothness while satisfying the orthogonality constraints (28) to within a reasonable epsilon (less than $1\text{e-}3$). This amounts to the nonlinearly constrained optimization problem of finding the solution which minimizes the deviation from orthogonality while maximizing the Sobolev exponent.

This approach led to a new family of QMFs; the 7- and 9-tap examples are given in Table 2 below. The 9-tap Sobolev-optimal QMF is quite distinct from the 9-tap QMF of Adelson-Simoncelli[8]; its frequency response is plotted in Figure 11, while pictures comparing the corresponding scaling and wavelet functions are shown in Figure 12. Observe that our Sobolev-optimal design yielded a flatter frequency response than the previous frequency-sampling approach [8]. The scaling function associated with the Sobolev-optimal QMF has Sobolev exponent $s = 1.795$, and is visibly smoother

than the scaling function associated with the Adelson-Simoncelli QMF, which has $s = 1.567$. Both of these scaling functions are smoother than the analysis scaling function of the Daubechies (7,9)-tap pair, which has $s = 1.410$. The Sobolev-optimal 9-tap QMF has a maximum deviation from orthogonality of $6.6e - 04$, superior to the $7.9e - 04$ max deviation for the Adelson-Simoncelli 9-tap QMF. The new 7-tap Sobolev-optimal QMF proved to be quite similar to Adelson and Simoncelli's 7-tap design.

n	$a_0[n]$	n	$a_0[n]$
0	-7.8125000e-03	0	1.9531250e-02
1	-7.3069675e-02	1	-4.6875000e-02
2	3.6136589e-01	2	-7.3234012e-02
3	8.5324613e-01	3	4.0042839e-01
		4	8.1451231e-01

Table 2: 7 and 9-tap Sobolev-optimal QMF lowpass filters.

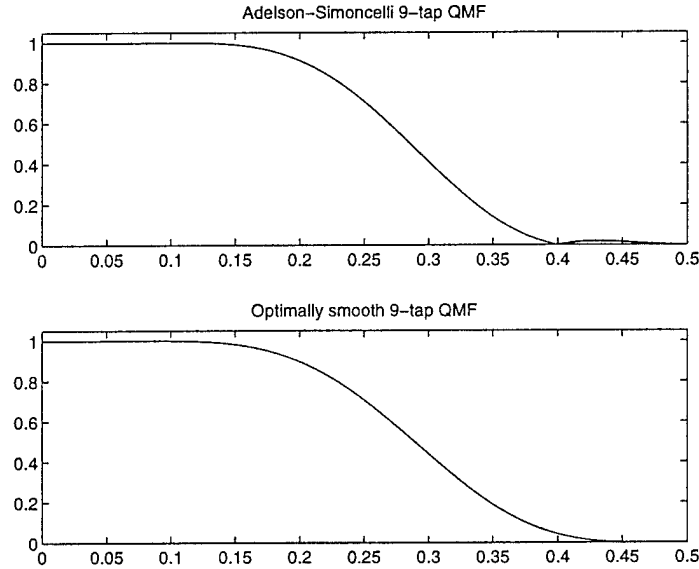


Figure 11: Frequency responses of Adelson-Simoncelli and optimally smooth lowpass filters.

We applied the new QMF banks to lossy wavelet-based image compression of natural and fingerprint images, and found the 9-tap QMF to be competitive with the Adelson-Simoncelli QMF in both numerical and qualitative senses; both of these wavelet filters were superior to the Daubechies (7,9)-tap pair at compressing a fingerprint image with the FBI's WSQ tree. Further details are given in [48].

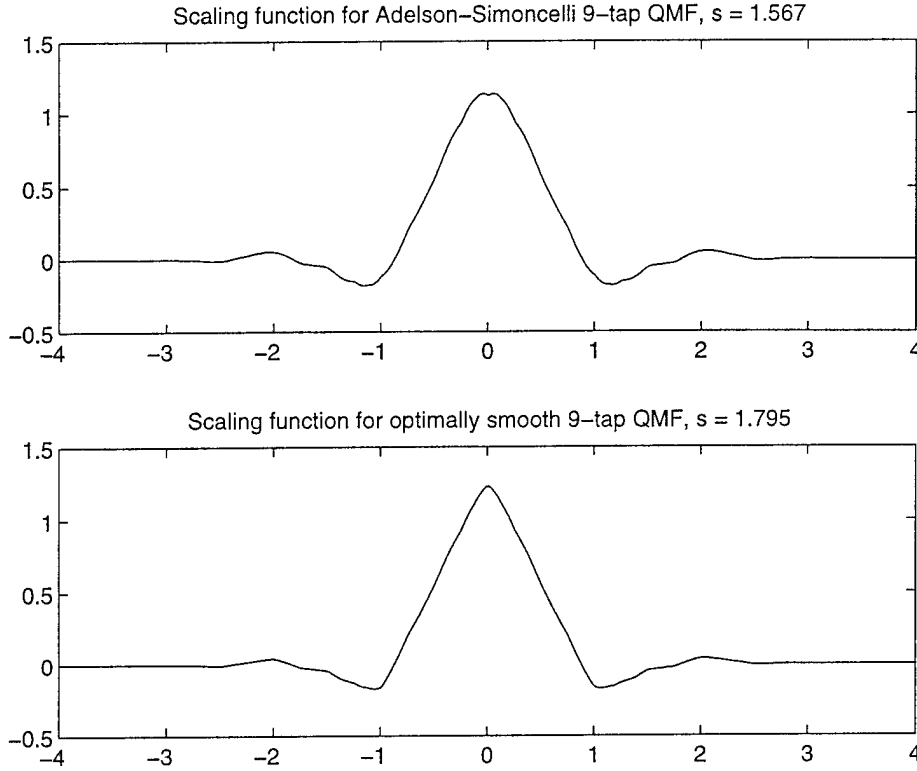


Figure 12: Comparison of Adelson-Simoncelli and optimally smooth scaling functions.

4.4.5 Maximally Hölder-smooth Wavelet Filters

The maximal-smoothness approach was also applied to the design of 2-channel orthogonal wavelets, but this time using the time-domain notion of Hölder regularity, as opposed to the L^2 -frequency domain notion of Sobolev regularity. We now briefly summarize our approach and the results; further details appear in [58].

Our method utilized the autocorrelation sequence domain where 1) the dimension of the parameter space can be considerably reduced and 2) an *unconstrained* optimization problem results (as opposed to a *constrained* problem if one were to work in the impulse response domain). We maximized the Hölder smoothness over the class of wavelet autocorrelation sequences of a given length.

Hölder smoothness is a direct generalization of the notion of differentiability from the natural to the real numbers. A function f has Hölder smoothness of order $\alpha \in [0, 1)$ if, for any t, h ,

$$|f(t+h) - f(t)| < c|h|^\alpha, \quad (32)$$

where $c < \infty$ is a constant independent of t and h . For a higher order of Hölder smoothness $s_H(f) = L + \alpha$, $L \in \mathbb{N}$, $\alpha \in [0, 1)$, the definition above is applied to

the L -th derivative of f . It can be shown that a function f possesses L continuous derivatives if and only if it has Hölder smoothness of order $s_H(f) \geq L$. Rioul has given an algorithm for determining the Hölder smoothness [86] that can be easily implemented and yields a very accurate estimate for s_H . Given the definition of Hölder smoothness, the optimization problem to be solved can be stated as follows:

Maximize $s_H(\varphi)$ over the coefficients $\mathbf{p} = [p_K, p_{K-2}, \dots, p_1]^T$ of $P(\omega)$
according to equations (2)-(3), and subject to

$$P(\omega) \geq 0. \quad (33)$$

In order to compute the scaling filter (and not its autocorrelation sequence), one must find the spectral factor [24] of the solution \mathbf{p} .

The solution employs a general purpose optimization algorithm requiring only a function that computes the order of smoothness s_H for any valid set of parameters. The approach outlined here works as well for any other method of measuring smoothness, such as Sobolev differentiability ([31], [48], [120]); and further explorations are described in [59].

It is obviously desirable to reduce the number of free parameters to reduce the complexity of the problem. One can do so by means of the previously quoted theorem relating wavelet differentiability to vanishing moments [24]. For a wavelet scaling function to have N continuous derivatives, the z -transform $P(z)$ must have $2(N+1)$ zeros at $z = -1$. Thus the dimension of the parameter space can be reduced to $\frac{K-1}{2} - N$. Although we do not know the achievable order of smoothness for a given filter length $K+1$, we have a lower bound by measuring the Hölder smoothness of the corresponding Daubechies filter. For example, the D6 filter has 3 free parameters but $s_H \approx 1.08$, so that there is just $3 - N - 1 = 1$ parameter left for optimization.

The remaining free parameters are further constrained by the nonnegativity condition (33). It effectively constrains the possible root locations of $P(z)$. Single roots on the unit circle are not admissible; complex roots off the unit circle have to occur in quadruples, real roots must occur in mirror pairs, and roots on the unit circle must come in complex conjugate pairs with multiplicity two, thus fixing two, one, and two parameters, respectively. It is important to note that each of the three possible cases can be expressed by equations that are linear in the autocorrelation coefficients.

However, it is not clear *a priori* which of the different combinations of roots (e.g., in the case of two free parameters either one complex quadruple or a pair of double roots on the unit circle can be chosen) leads to the maximum smoothness. Thus one has to try all possible constellations. Our algorithm may be summarized as follows:

- Choose the length $K+1$ of the scaling filter.
- Determine the number of degrees of freedom, i.e., find how many vanishing moments can be relaxed.

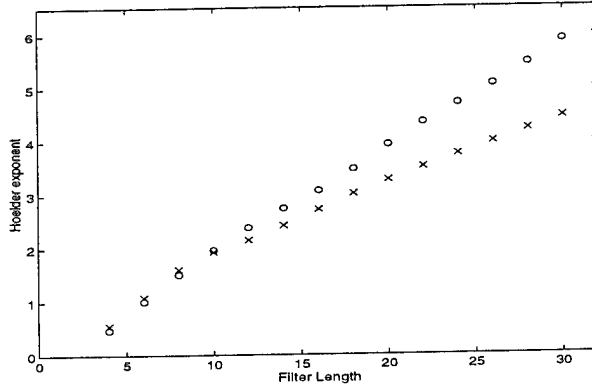


Figure 13: Hölder smoothness for Daubechies (\times) and optimized (\circ) filters as a function of filter length.

- Determine all possible constellations that could give rise to a nonnegative power spectrum.
- For each possible constellation, numerically maximize the desired smoothness over the parameter space, i.e., the $\frac{K-1}{2} - N$ radii and angles. A general purpose optimization routine is used. It only requires a function that computes the order of smoothness, given the parameter values.

We carried out the above method for filter lengths ranging from 4 to 30. When the filter length was 10 or larger, we obtained filters with greater Hölder smoothness than the Daubechies orthogonal wavelets. For lengths 4, 6, and 8, the optimization procedure did not yield higher smoothness than the Daubechies' filters. Figure 13 shows the the Hölder exponent s_H as a function of $K + 1$ for the scaling functions corresponding to Daubechies filters and to the new filters. Beyond length 8 the new method clearly improves the smoothness, with increasing success for larger filter lengths. A linear least-squares fit to each of the smoothness plots found the Daubechies filters with order less than 30 to have a Hölder exponent growing as $.1418(K + 1)$, while the maximally smooth filters have a Hölder exponent growth of $.2024(K + 1)$. We know [122] that asymptotically the Daubechies filters' Hölder exponents grow as $.1037(K + 1)$. The asymptotic behavior of the maximal possible Hölder exponent for a given filter length, as the length increases, remains an open question.

In Figure 14 we have plotted the number of zeros at π for the Daubechies and maximally smooth scaling filters; this is equal to the number of vanishing moments of the corresponding wavelet system. Notice that while Daubechies filters of length $K + 1$ have a zero of order $\frac{K+1}{2}$ at π , the maximally smooth filters have a zero of roughly half that order. In other words, we have given up half the vanishing moments of the Daubechies filter to obtain free parameters for optimizing the Hölder regularity. Notice that this behavior is not monotonic, probably because several different root

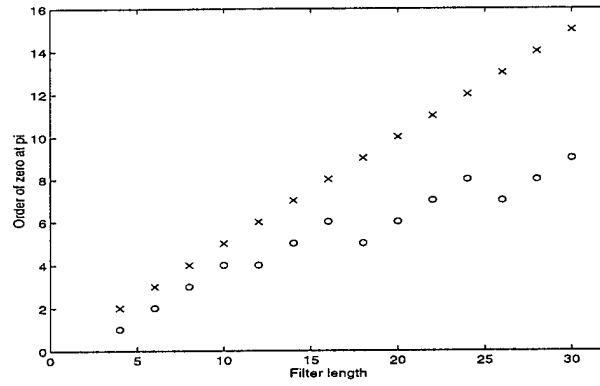


Figure 14: Order of zero at π (number of vanishing moments) for Daubechies (\times) and maximally smooth (\circ) filters as a function of filter length.

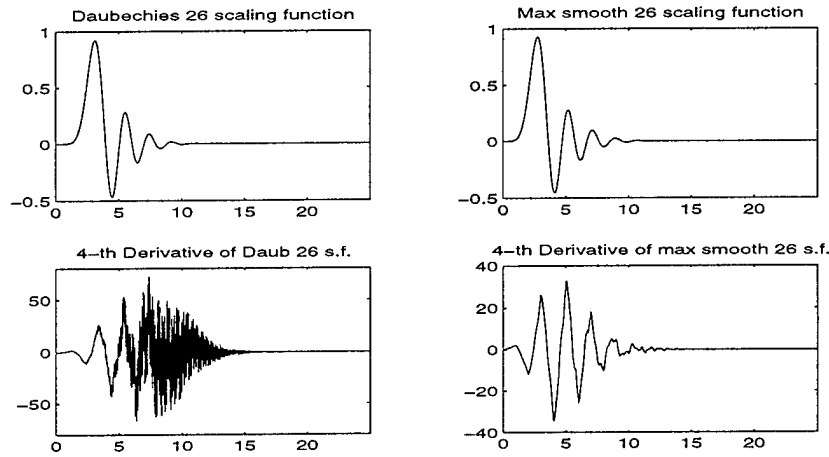


Figure 15: Scaling functions (top) and their 4-th derivatives (bottom) corresponding to 26-coefficient Daubechies and Hölder-optimized filters.

constellations lead to near-optimal filters.

Consider the case of a length 26 ($K = 25$) scaling filter that is optimized for s_H . The Daubechies filter corresponds to smoothness $s_H = 4.005$. After optimization of three distinct complex quadruples (off the unit circle) one gets a scaling function with smoothness $s_H = 5.06$, one more continuous derivative. Figure 15 depicts the scaling functions and their 4-th derivatives for each of these filters. While the scaling functions have a similar appearance, the higher degree of smoothness of the derivative is quite apparent. The filter frequency responses are shown in Figure 16.

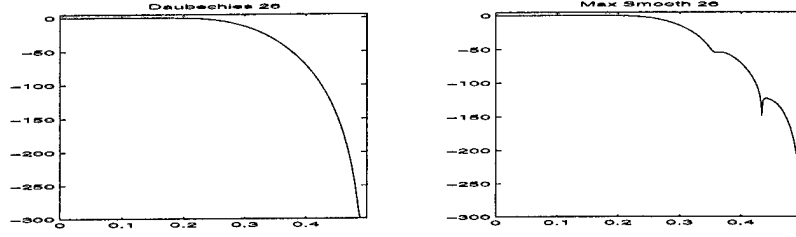


Figure 16: Frequency responses in dB of the 26-coefficient Daubechies (left) and Hölder-optimized (right) filters.

4.4.6 A Frequency-sampling Approach to Smooth Wavelet Filters

In addition to simply optimizing for maximally smooth filters, one may use the underlying mathematics of wavelet regularity to inform the design of wavelet lowpass filters, while using the traditional technique of frequency sampling. Frequency sampling simply means specifying the value of a filter response at particular points in the frequency domain; in this case the value specified will be zero.

Recall that the asymptotic regularity estimates of section 4.3 were expressed in terms of the values of the remainder trigonometric polynomial $R(\omega)$ at the periodic point ω_c of the ergodic map $\tau : \omega \rightarrow M\omega \bmod 2\pi$ and its preimages. Notice as well that the Daubechies and generalized Daubechies scaling filters are specified by zeros at the preperiodic points $\omega = 2\pi k/M$, $1 \leq k \leq M-1$ (when $M=2$, this reduces to an N -th order zero at π). Given these examples, we have been led to define a new class of scaling functions and corresponding wavelet systems that includes the maximal vanishing moments family as a special case. The new class consists of those wavelet scaling filters (i.e. sequences satisfying (2) and (3)) whose Fourier transform has specified zeros of specified orders at preperiodic points of the map τ . The lowpass filters are specified by a frequency sampling condition. We call these PPZ wavelet matrices (for PrePeriodic Zero) and they are determined by the choice of preperiodic zeros at specified locations ω_i of order N_i ; this collection of constraints is designated generically by $(\omega_1; N_1, \omega_2, N_2; \dots; \omega_L, N_L)$. Note that the sum condition (3) simply amounts to requiring a zero of order 1 at $2\pi k/M$, $1 \leq k \leq M-1$ for each PPZ wavelet matrix. We have investigated several PPZ wavelet matrix families for the cases $M=2, 3, 4$ and will show how they can be used to improve the Sobolev regularity of a filter with a given length.

For $M=2$, the dyadic wavelet case, we consider three PPZ wavelet matrix families with filter length $2N$ and zeros of the specified locations and orders:

I_N : (π, N) (Daubechies),

II_N : $(\pi, N-2; 2\pi/3, 1; 4\pi/3, 1)$

III_N : $(\pi, N-4; 2\pi/3, 1; 4\pi/3, 1; 5\pi/6, 1; 7\pi/6, 1)$

The definitions of families II_N and III_N were motivated by the role of $R(\frac{2\pi}{3})$ and

$R(\frac{5\pi}{6})$ in the asymptotic estimates of section 4.3. These two families can be constructed explicitly; instead of using the minimal length remainder $R_N(\omega)$ for N vanishing moments, use the more general expression

$$R(\omega) = R_N(\omega) + (1 - \cos \omega)^N \sum_{n \neq Mk} \tilde{r}_n \cos n\omega, \quad (34)$$

and choose the free parameters \tilde{r}_n to assert additional zeros at the points ω_i . For example, the representative for family II_N with filter length $2N$ will be determined by

$$R(\omega) = R_{N-1}(\omega) + (1 - \cos \omega)^{N-1} (\tilde{r}_1 \cos \omega + \tilde{r}_3 \cos 3\omega),$$

where \tilde{r}_1 and \tilde{r}_3 are determined by the conditions

$$R\left(\frac{2\pi}{3}\right) = 0; \quad R'\left(\frac{2\pi}{3}\right) = 0.$$

$R(\omega)$ must have a double zero at $\omega_c = \frac{2\pi}{3}$ because it is associated with the magnitude squared of the Fourier transform. The coefficients of $R_N(\omega)$ are predetermined by formulae (16)-(18), and can be explicitly solved for.

Both families II_N and III_N have Sobolev regularity growing faster with N than the $.2075N$ growth of the Daubechies orthogonal wavelets, and the Sobolev regularity of the family III_N grows faster than that of II_N , as shown in Figure 17. Notice that family II_N does not yield a smoother scaling function until the sequence length reaches 20, and family III_N does not do so until the sequence length reaches 18. Figure 18 shows the scaling functions and their fourth derivatives for representatives of each of these families corresponding to length-30 scaling sequences. The Sobolev exponents of the three scaling functions are 4.565, 5.410, and 6.291, respectively. While the scaling functions appear quite similar, the difference shows up in the graph of the fourth derivative.

Similarly, when $M = 4$ we examined two PPZ wavelet families, the maximal vanishing moment family with filter length $4N$

$$I_N : (\pi, N; \pi/2, N; 3\pi/2, N) \text{ (generalized Daubechies),}$$

and a new family with filter length $4N - 1$

$$II_N : (\pi, N - 1; \pi/2, N - 1; 3\pi/2, N - 1; 4\pi/5, 1; 6\pi/5, 1).$$

The new PPZ family has one vanishing moment relaxed, with the resulting degrees of freedom used to assert a zero at the periodic point $4\pi/5$. The scaling functions from this family proved to have Sobolev regularity that grew significantly faster than for the maximal vanishing moment family I_N , as shown in Figure 19.

For $M = 3$ we define the families:

$$I_N : (2\pi/3, N; 4\pi/3, N) \text{ (generalized Daubechies, sequence length } 3N)$$

$$II_N : (2\pi/3, N - 1; 4\pi/3, N - 1; \pi, 1) \text{ (sequence length } 3N - 1)$$

and find that the asymptotic smoothness of family I_N is logarithmic in N (as mentioned above), and algebraic in N (of order $N^{.55}$) for family II_N . Figure 20 shows the

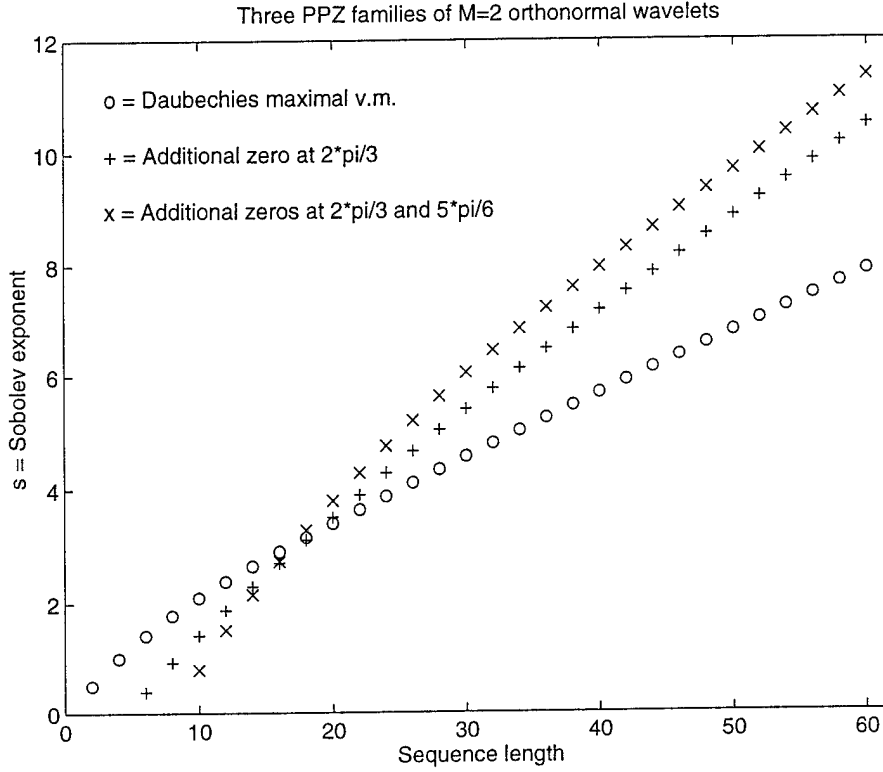


Figure 17: Sobolev exponents for the PPZ families I_N , II_N , and III_N of $M = 2$ orthonormal wavelets.

Sobolev exponents for scaling functions from the new family and the minimal-length family. Scaling functions from the two $M = 3$ families with similar support lengths (17 and 17.5) are compared in Figure 21.

These examples illustrate the relation between smoothness and the specification of the Fourier transform of the filters at specified preperiodic points, as pointed out by previous authors [20], [24]. The patterns described indicate that a family of PPZ wavelets with vanishing specified at preperiodic points other than just the preimages of 2π under τ may, in general, lead to greater smoothness than the maximal vanishing moment family.

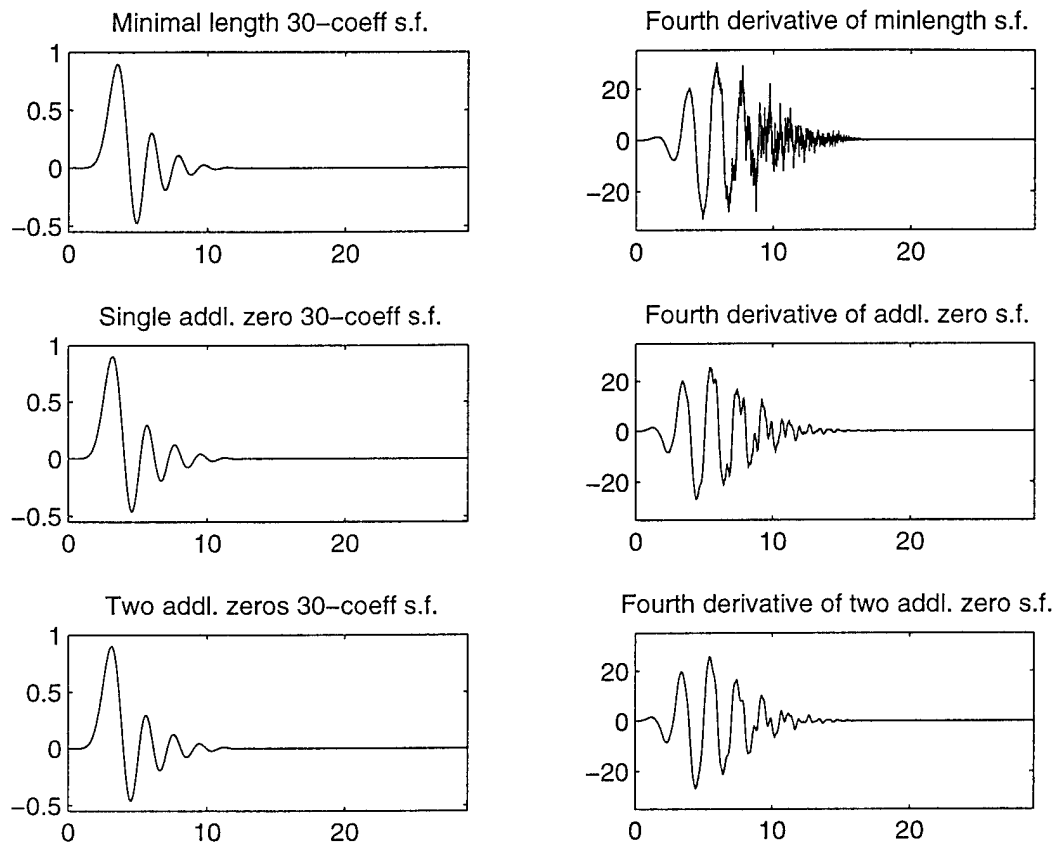


Figure 18: Dyadic ($M = 2$) scaling functions and their fourth derivatives for 30-coefficient representatives of the families I_N , II_N , and III_N . From top to bottom: Daubechies minimum length, single additional zero at $2\pi/3$, and two additional zeros at $2\pi/3$ and $5\pi/6$.

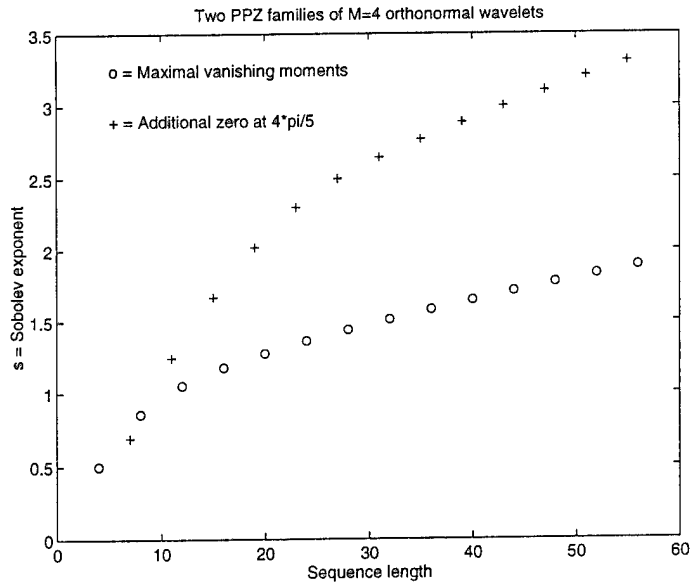


Figure 19: Sobolev exponents for the families I_N and II_N of $M = 4$ wavelets.

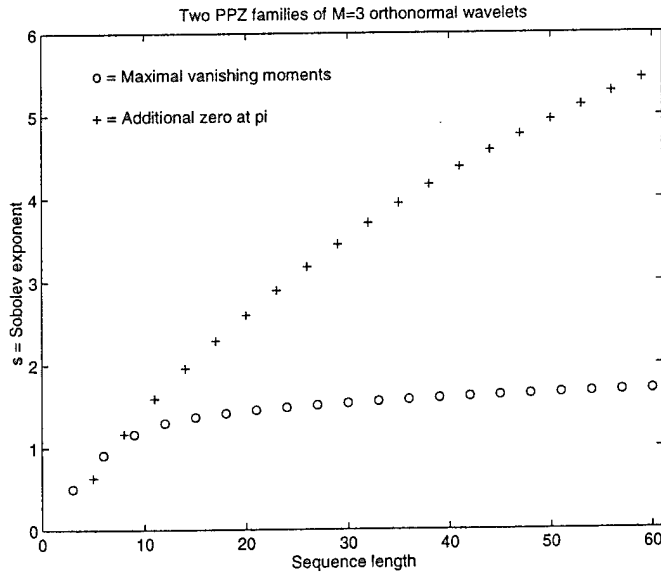


Figure 20: Sobolev exponents for the two families I_N and II_N of $M = 3$ scaling functions, $N = 1, 2, \dots, 15$.

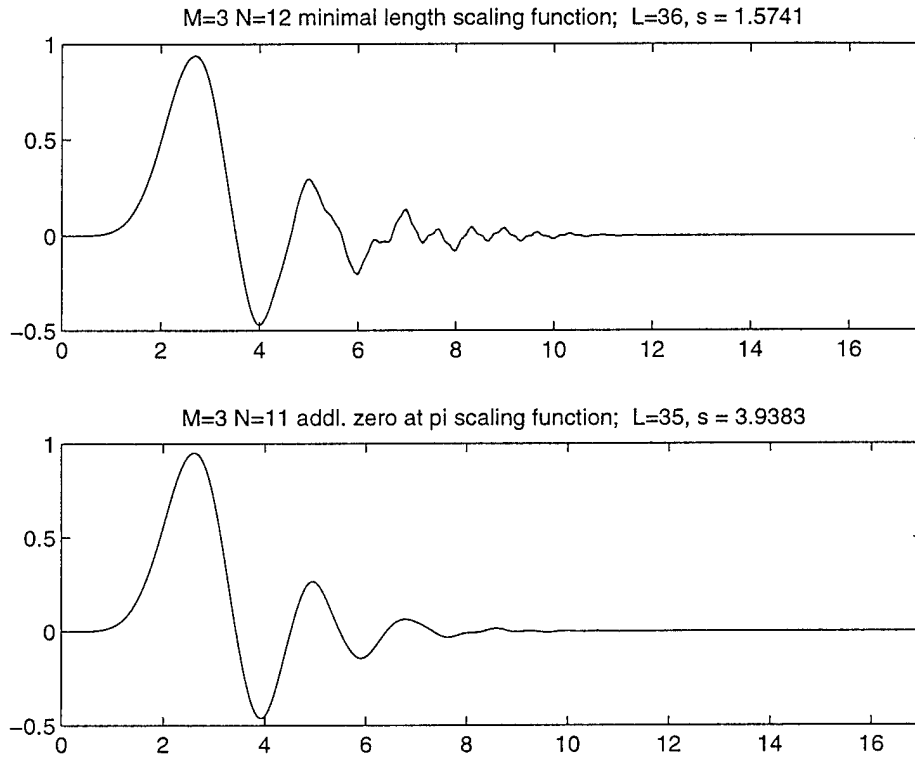


Figure 21: Two $M = 3$ scaling functions, one with support 18 and maximal approximation degree (12), and the other with support $17\frac{1}{2}$, approximation degree 11, and an additional zero at π for the symbol.

4.4.7 Quadratic-Constrained Optimization for Wavelet Filter Design

In addition to developing the approaches to wavelet filter design described above, we have implemented a high-performance design approach due to Professor Truong Q. Nguyen of the University of Wisconsin, a consultant to Aware on this contract. Nguyen's QCLS (quadratic-constrained least-squares) method [69], [70] is applicable to a broad class of filter design problems, including a wide variety of wavelet filter banks. This approach formulates the orthogonality condition (2) as a quadratic constraint on the wavelet filter coefficients, and optimizes a functional such as out-of-band energy over the coefficient space (i.e. using the time-domain representation of the filter). It is possible to include vanishing wavelet moments as additional quadratic constraints. QCLS has produced the highest-performance (best stopband attenuation or subchannelization) design examples known for cosine-modulated filter banks (cosine packets). We describe this method briefly now.

The QCLS algorithm begins with the vector \mathbf{h} of unknown lowpass filter coefficients, and formulates the perfect-reconstruction or orthogonality conditions as a quadratic form in \mathbf{h} . Depending on the structure of the rank M wavelet matrix (arbitrary paraunitary matrix, linear-phase, or cosine-modulation, as described in section 4.5 below), the number N_c of constraints and their particular form will vary. In all cases, the N_c independent PR constraints in (2) can be expressed as

$$\mathbf{h}^t \mathbf{Q}_\ell \mathbf{h} = c_\ell, \ell = 1, 2, \dots, N_c.$$

The objective function Φ for the optimization (taken to be the L^2 stopband error of $H(z)$) can also be expressed as a quadratic form in \mathbf{h} via the eigenfilter formulation [114]:

$$\Phi(\mathbf{h}) = \int_{\omega_s}^{\pi} |H(e^{j\omega})|^2 d\omega = \mathbf{h}^t \mathbf{P} \mathbf{h}.$$

Both the objective function and constraints admit closed-form expressions for the gradient and Hessian, enabling efficient solution of the constrained minimization problem

$$\min_{\mathbf{h}} \Phi(\mathbf{h}) \quad \text{subject to} \quad \mathbf{h}^t \mathbf{Q}_\ell \mathbf{h} = c_\ell, 1 \leq \ell \leq N_c. \quad (35)$$

A nonlinear constrained-optimization algorithm such as that of Schittkowski [91] is used to solve (35). The solution \mathbf{h} is the time-domain impulse response of a low-pass filter with minimal stopband energy, satisfying the appropriate rank M wavelet orthogonality conditions.

We have implemented the QCLS approach for the design of cosine-modulated rank M wavelet matrices in our WaveTool software, as well as the earlier lattice-parametrization method which uses an unconstrained optimization over a highly nonlinear parameter space. The QCLS method systematically yields higher-performance filters, as shown in the two WaveTool plots below (Figures 22 and 23).

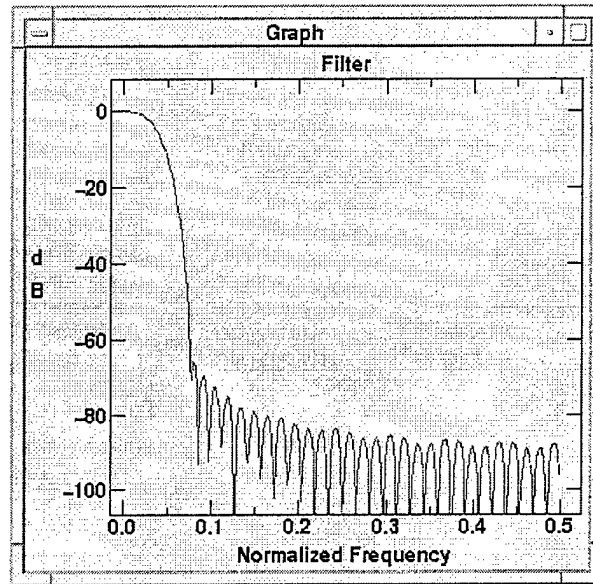


Figure 22: Frequency response of an $M = 8$ length 64 lowpass filter for cosine modulation, designed with QCLS using WaveTool software.

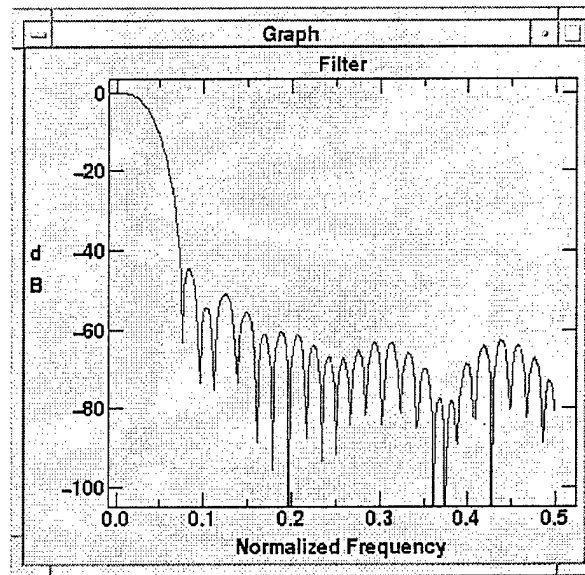


Figure 23: Frequency response of an $M = 8$ length 64 lowpass filter for cosine modulation, designed with lattice parametrization using WaveTool software.

4.5 Construction of Full Rank M Wavelet Matrices

In the previous section we saw a number of methods for constructing rank M scaling sequences of degree N and arbitrary length. We now turn to the design of the corresponding wavelet sequences, i.e. to the problem of constructing a complete rank M wavelet matrix given its first row. In the rank M case, there is considerable freedom in such a construction. We describe here three distinct methods. The first is completely general, and is based on the prior choice of a wavelet lowpass filter (scaling sequence) and a Haar wavelet matrix [46]. The second and third methods involve more specialized structures – linear phase rank M wavelet matrices and rank M cosine-modulated wavelets, respectively. All three approaches have associated fast algorithms for computation.

4.5.1 General Haar-based constructions

First, we present a method [39] for constructing a full wavelet matrix given its first row (the scaling sequence) and an $M \times M$ matrix which we call the characteristic Haar matrix [46].

A Haar wavelet matrix is an orthogonal matrix (up to scalar multiplication) whose first row is all ones; that is, its entries $h_{s,k}$ satisfy

$$\sum_k h_{s,k} h_{s',k} = M \delta_{s,s'} ,$$

and $h_{0,k} = 1 \ \forall k$.

Observe that every such Haar wavelet matrix is a rank M wavelet matrix (with $K = M$) and that every $M \times M$ wavelet matrix is a Haar matrix. Useful examples of Haar wavelet matrices include the M -point FFT, type II Discrete Cosine Transform, and Hadamard matrix. The collection of rank M Haar matrices is isomorphic to the group of orthogonal matrices of rank $M - 1$.

It will serve us to think of our wavelet matrices as being $M \times Mg$ for some integer g ; we can always pad each row with zeros to bring the wavelet matrix into this form. g is called the *genus* of the wavelet matrix. If we break up the $M \times Mg$ wavelet matrix \mathbf{A} into its constituent $M \times M$ blocks

$$\mathbf{A} = (\mathbf{A}_0 \ \mathbf{A}_1 \ \dots \ \mathbf{A}_{g-1}) \tag{36}$$

then the *characteristic Haar matrix associated with \mathbf{A}* is given by

$$\mathbf{H}_0 = \mathbf{A}_0 + \mathbf{A}_1 + \dots + \mathbf{A}_{g-1} .$$

It can be checked that \mathbf{H}_0 is in fact a Haar wavelet matrix.

In [46] we solved the following problem: given a Haar wavelet matrix \mathbf{H}_0 and a scaling sequence a_0 , construct a full wavelet matrix \mathbf{A} whose first row is a_0 and

whose characteristic Haar matrix is \mathbf{H}_0 . That explicit construction can be clarified and refined using Vaidyanathan's paraunitary factorization technique. In particular, a parametrization of the choice of the $M - 1$ wavelets is given by the choice of the characteristic Haar matrix.

Working in the z -transform domain, Vaidyanathan [116] has proven that *every* paraunitary polyphase matrix $\mathbf{H}(z)$ of McMillan degree¹ K can be factored into the form

$$\mathbf{H}(z) = \left(\prod_{k=0}^{K-1} (\mathbf{I} - \mathbf{v}_k \mathbf{v}_k^\dagger + z \mathbf{v}_k \mathbf{v}_k^\dagger) \right) \mathbf{H}_0, \quad (37)$$

where each \mathbf{v}_k is a unit M -vector. $\mathbf{H}(z)$ will be the polyphase matrix of a wavelet matrix if and only if \mathbf{H}_0 is a Haar wavelet matrix, so (37) provides a factorization of all wavelet matrices with polyphase matrix of McMillan degree K . We refer to the term

$$\mathbf{I} - \mathbf{v}_k \mathbf{v}_k^\dagger + z \mathbf{v}_k \mathbf{v}_k^\dagger$$

as a *prime factor* of the polyphase matrix.

Theorem 4.6 *Given a scaling sequence a_0 of overlap g and a characteristic Haar matrix \mathbf{H}_0 , there exists a unique wavelet matrix of McMillan degree $g - 1$ (i.e. with $g - 1$ prime factors) whose first row is a_0 and with characteristic Haar \mathbf{H}_0 . Furthermore, we can explicitly construct the prime factors $\mathbf{I} - \mathbf{v}_k \mathbf{v}_k^\dagger + z \mathbf{v}_k \mathbf{v}_k^\dagger$ (and thus the wavelet matrix) from a_0 and \mathbf{H}_0 .*

The proof of this theorem (including explicit construction of the prime factors) appears in [39].

As an example, we use this method to construct a wavelet matrix with $M = 4$ and $g = 2$, with approximation degree $N = g = 2$. The minimal length lowpass filter (scaling sequence) for this case was given previously (equation (20)). The full wavelet matrix constructed using Theorem 4.6, with the sequence (20) for its first row and the rank-4 DCT-II for its characteristic Haar matrix \mathbf{H}_0 is

$$\begin{pmatrix} 0.5396 & 0.7896 & 1.0396 & 1.2896 & 0.4604 & 0.2104 & -0.0396 & -0.2896 \\ -0.1962 & -0.1456 & -0.4120 & -0.3614 & 1.5028 & 0.6868 & -0.1292 & -0.9451 \\ 1.0 & -1.0 & -1.0 & 1.0 & 0.0 & 0.0 & 0.0 & 0.0 \\ 0.4344 & -1.3554 & 1.3157 & -0.4740 & 0.1068 & 0.0488 & -0.0092 & -0.0672 \end{pmatrix}.$$

¹A polyphase matrix of McMillan degree K will correspond to a wavelet matrix of overlap $K + 1$, while a wavelet matrix of overlap $K + 1$ has a polyphase matrix with McMillan degree at least K . However, there exist wavelet matrices of overlap $K + 1$ and McMillan degree strictly greater than K ; for examples see [50]. The construction summarized here and presented in detail in [39] describes a unique wavelet matrix with first row a_0 and characteristic Haar \mathbf{H}_0 and having a polyphase matrix of McMillan degree K .

The frequency responses of the lowpass filter (scaling sequence) and the three wavelet filters (i.e., the four rows of the matrix) appear in Figure 24.

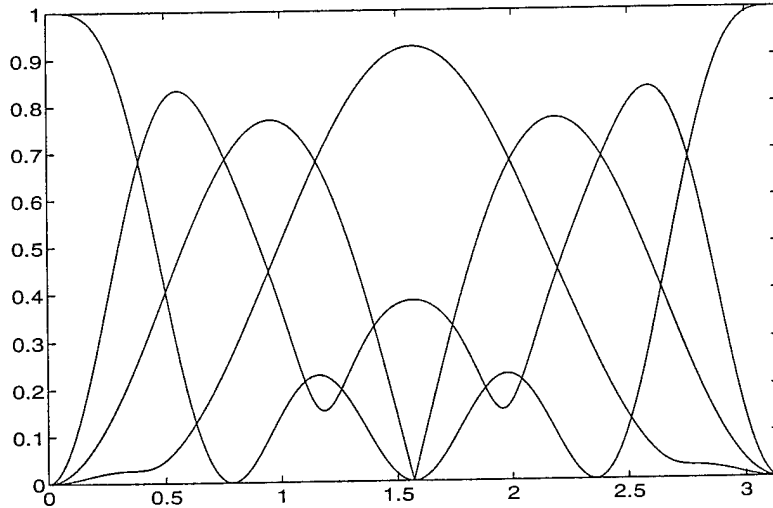


Figure 24: Magnitude frequency responses of filters in a minimal length $M = 4$, $N = 2$ wavelet matrix based on a *DCT-II* characteristic Haar matrix.

When the rank $M > 2$, the construction of Theorem 4.6 gives considerable freedom in the choice of the $M - 1$ wavelet sequences. We have exploited the range of possibilities in this parametrization for image compression applications, as reported in [45].

4.5.2 Linear-Phase Rank M Wavelets, or GenLOT constructions

A more restrictive subclass of the rank M wavelets are those with linear-phase, i.e. such that each of the M filters have linear-phase symmetry. Such symmetry is of great utility in applications to image compression, where symmetric extension of data is the preferred method for handling image boundaries [95]. Linear-phase filters are also of value in a wavelet transform because they preserve centers of mass in an iterated signal decomposition. In the rank 2 case, wavelet matrices cannot simultaneously satisfy orthogonality (2) and symmetry, other than the nearly trivial Haar example. This has been overcome by designing biorthogonal rank 2 wavelet matrices [73], [11], [117]. However, when $M > 2$, linear-phase and orthogonality can be simultaneously satisfied. Such orthogonal linear-phase filter banks have recently been parametrized [97], [79], at least when M is even. We have used these parametrizations to determine the M -band orthogonal linear-phase wavelets, and successfully applied the new

wavelet matrices to problems in image compression.

Let us briefly summarize the parametrizations of Soman and deQueiroz. One of these works [79] identifies M -band orthonormal filter banks in which each filter has linear-phase symmetry as generalizations of the lapped orthogonal transform (LOT) [66], and confers the name GenLOT. As described previously, a rank M wavelet matrix can be described in the z -transform domain by its polyphase matrix $\mathbf{E}(z)$:

$$\begin{bmatrix} \mathbf{H}_0(z) \\ \mathbf{H}_1(z) \\ \vdots \\ \mathbf{H}_{M-1}(z) \end{bmatrix} = \mathbf{E}(z) \begin{bmatrix} 1 \\ z^{-1} \\ \vdots \\ z^{-(M-1)} \end{bmatrix}.$$

When the rank M is even, and the filters $h_k[n]$ (with z -transforms $\mathbf{H}_k(z)$) have length MN , the polyphase matrix has the form

$$\mathbf{E}(z) = \mathbf{K}_{N-1}(z)\mathbf{K}_{N-2}(z)\dots\mathbf{K}_1(z)\mathbf{E}_0 \quad (38)$$

where each

$$\mathbf{K}_i(z) = \frac{1}{2} \begin{bmatrix} \mathbf{U}_i & \mathbf{0} \\ \mathbf{0} & \mathbf{V}_i \end{bmatrix} \begin{bmatrix} \mathbf{I} & \mathbf{I} \\ \mathbf{I} & -\mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{0} & z^{-1}\mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{I} & \mathbf{I} \\ \mathbf{I} & -\mathbf{I} \end{bmatrix}.$$

\mathbf{I} is the rank $M/2$ identity matrix, while \mathbf{U}_i and \mathbf{V}_i are arbitrary rank $M/2$ orthogonal matrices. \mathbf{E}_0 is a rank M unitary matrix with $M/2$ symmetric and $M/2$ antisymmetric rows (such as the DCT-IV); it can be factored as

$$\mathbf{E}_0 = \frac{1}{\sqrt{2}} \begin{bmatrix} \mathbf{D}_0 & \mathbf{0} \\ \mathbf{0} & \mathbf{D}_1 \end{bmatrix} \begin{bmatrix} \mathbf{I} & \mathbf{I} \\ \mathbf{I} & -\mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{0} & \mathbf{J} \end{bmatrix};$$

\mathbf{J} is the familiar reverse identity matrix (of rank $M/2$), while \mathbf{D}_0 and \mathbf{D}_1 are arbitrary orthogonal matrices of rank $M/2$.

This GenLOT formulation can be used to construct of M -band linear-phase orthonormal wavelets. First, all possible GenLOTs with one vanishing moment may be described in terms of certain rotation matrices. As described in Theorem 4.1, an orthogonal filter bank has one vanishing moment if it satisfies

$$\begin{bmatrix} \mathbf{H}_0(z) \\ \mathbf{H}_1(z) \\ \vdots \\ \mathbf{H}_{M-1}(z) \end{bmatrix} = \begin{bmatrix} \sqrt{M} \\ 0 \\ \vdots \\ 0 \end{bmatrix} \quad (39)$$

for $z = 1$. This is equivalent to the linear equation (3) in the definition of a wavelet matrix. We find that a GenLOT with $N - 1$ factors $\mathbf{K}_i(z)$ will have one vanishing

moment if and only if the rotation matrices \mathbf{U}_i satisfy

$$\mathbf{U}_{N-1}\mathbf{U}_{N-2}\dots\mathbf{U}_1\mathbf{D}_0 \begin{bmatrix} 1 \\ 1 \\ \vdots \\ 1 \end{bmatrix} = \begin{bmatrix} \sqrt{M/2} \\ 0 \\ \vdots \\ 0 \end{bmatrix}.$$

Clearly, the number of free parameters available for filter design (to optimize such properties as stopband attenuation or coding gain) increases with both the number of channels M and the filter length, which is determined by the number of factors N . With even the simplest $N = 1$ LOTs, it is possible to create a filter bank based on the DCT-IV having one vanishing moment for use in a wavelet decomposition. A 4-band GenLOT with $N = 2$ (filter length 12) and one vanishing moment is shown in Figure 25.

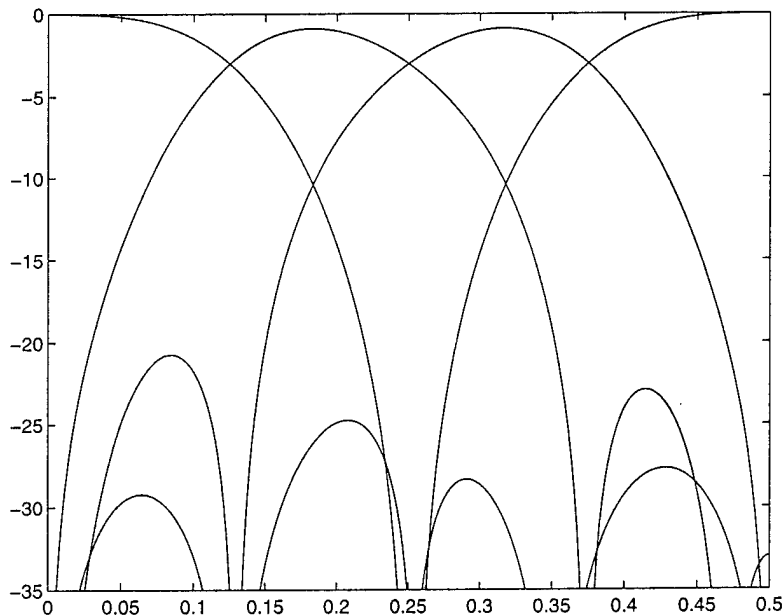


Figure 25: Magnitude responses (in dB) of a rank 4 genus 3 linear-phase wavelet matrix (GenLOT) with one vanishing moment.

It is of interest to create wavelet filters with more than one vanishing moment; indeed, the interpolation/approximation properties of such filters lead to their superiority for wavelet-based image coding systems [63]. An rank M wavelet matrix will

have two vanishing moments if it satisfies (39) as well as the second-order condition

$$\frac{d}{dz} \begin{bmatrix} \mathbf{H}_0(z) \\ \mathbf{H}_1(z) \\ \vdots \\ \mathbf{H}_{M-1}(z) \end{bmatrix} = \begin{bmatrix} * \\ 0 \\ \vdots \\ 0 \end{bmatrix} \quad (40)$$

at $z = 1$. In the GenLOT case, this leads to a set of linear constraint equations on the rotation matrices \mathbf{U}_i , \mathbf{V}_i , and \mathbf{D}_i that parametrize the wavelet matrix. These constraints can be solved to reduce the parameter space, enabling one to optimize over the remaining parameters to design a rank M linear-phase wavelet matrix with two vanishing moments and other desirable properties such as high stopband attenuation or maximal smoothness. The frequency responses of such a wavelet matrix with $M = 4$ are shown in Figure 26.

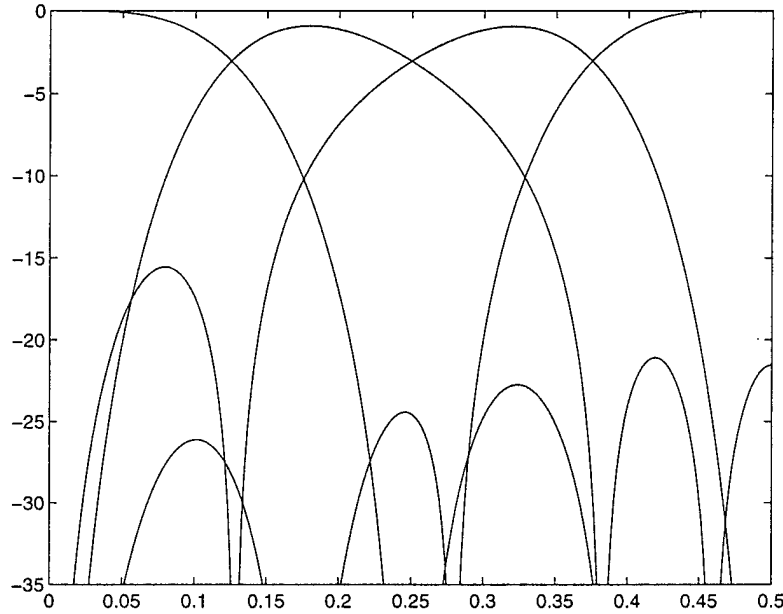


Figure 26: Magnitude responses (in dB) of a rank 4 genus 4 linear-phase wavelet matrix (GenLOT) with two vanishing moments and maximal Sobolev smoothness.

We applied these new higher rank, linear-phase wavelets to image compression [44], with results superior to those obtained using rank 2 wavelets. Particular advantages were found in the compression of fingerprints and seismic data. Rank 4 wavelet matrices appear naturally in the special wavelet-packet tree (Figure 27) specified by the FBI for their Wavelet Scalar Quantization fingerprint compression algorithm [7]; this transform tree can be obtained as the cascade of two rank 4 wavelet matrices, followed by a rank 2 wavelet in the lowest-frequency subband. When we substituted

the GenLOT examples given above into the transform structure, lower maximum (Chebyshev) errors resulted across all compression ratios; the effect was particularly pronounced at high compression ratios (low bitrates).

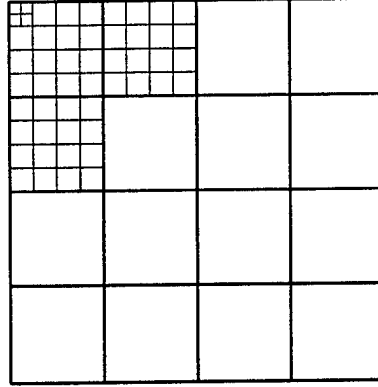


Figure 27: Wavelet transform tree specified by FBI for fingerprint compression.

4.5.3 Cosine-modulation

One of the most effective methods for the design and implementation of rank M wavelet filter bank transforms is via cosine modulation. In this approach, a single lowpass “prototype filter” is modulated to create a complete rank M wavelet matrix, using a DCT-based modulation matrix. This method yields both the highest-performance wavelet filters, and the fastest known algorithms for rank M wavelet transform computation. We summarize the cosine modulation construction here; such wavelet matrices also go by the name of cosine packets and discrete local cosine transforms.

In the most general biorthogonal case with DCT-IV modulation, the impulse responses of the analysis filters $h_k[n]$ are cosine-modulated versions of a prototype filter $h[n]$ of length N_h , and the synthesis filters can also be obtained via cosine-modulation of a length N_f prototype filter $f[n]$. The overall reconstruction delay D of the filter bank can be fixed arbitrarily in the range $D \in [0, N_f + N_h - 1]$. For a given delay $D = 2sM + d$ (where $0 \leq d < 2M$), the relation between the analysis filters, the synthesis filters, and their prototypes can be stated as follows:

$$h_k[n] = 2h[n] \cos \left[(2k+1) \frac{\pi}{2M} \left(n - \frac{D}{2} \right) + \theta_k \right], \quad (41)$$

$$f_k[n] = 2f[n] \cos \left[(2k+1) \frac{\pi}{2M} \left(n - \frac{D}{2} \right) - \theta_k \right], \quad (42)$$

with $\theta_k = (-1)^k \frac{\pi}{4}$. Note that the modulation does not depend on the filter length but only on the delay of the system. However, if we restrict ourselves to the case where $h[n] = f[n]$ and both are length N linear-phase prototype filters, the delay is constrained to be $D = N - 1$ and this is the same modulation as in [57].

As an example, the lowpass prototype filter shown in Figure 22 can be modulated to yield the rank 8 wavelet matrix (filter bank) shown in Figure 28 below. Notice that the superior stopband attenuation of the prototype filter is preserved in the modulated wavelet matrix. The superior subchannelization properties of cosine-modulation is due to a combination of the relatively clean parametrization of the class of possible prototype filters and the high-performance QCLS algorithm.

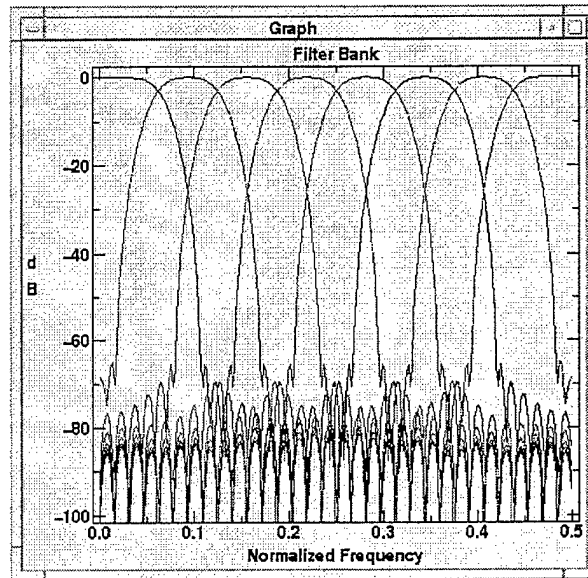


Figure 28: Frequency responses of an $M = 8$ length 64 cosine-modulated wavelet matrix, designed using WaveTool software.

Further details on cosine-modulated filter banks and wavelets appear in a number of references, including [66], [57], [81], [102], [36] [42] among others.

5 WaveTool Software

Foremost among our accomplishments under this contract was the development and release of the WaveTool software. This is a UNIX-based software tool for rapidly prototyping and evaluating general higher-rank wavelet decompositions, with a sophisticated graphical user interface. The WaveTool grew out of our early experience of applying wavelet methods here at Aware. We used to laboriously hand-code routines to perform wavelet/multirate decompositions for each new application as it arose. We had one set of software which did rank 2 wavelet decompositions for image compression, another which performed a rank 32 modulated lapped transform for CD-quality audio compression, and a third for transient detection. While each new application might require a new basis (new filters or new tree), the underlying mechanics of the decomposition was always the same. Furthermore, a significant step in the development of wavelet algorithms proved to be the hunt for the right basis. Out of this experience grew our specification for the WaveTool – a piece of software which would enable the user to rapidly prototype wavelet algorithms with a wide range of filter and tree choices, and underlying fast computer programs for implementing the wavelet decompositions, once chosen.

The WaveTool software provides the capability to design a variety of wavelet and multirate filter banks, assemble them into arbitrary tree structures, and process data through the system in either subband (one-signal-into-many-subbands) or transmultiplexer (many-signals-into-one) modes. The WaveTool also offers extensive graphical display capabilities for both filter bank components and input and output data. Perhaps the easiest way to describe the software is to take the reader on a tour of its functionality.

The heart of the WaveTool is its Tree Window, in which the user interactively defines a tree structure of paraunitary filter banks. He or she does so by picking off a menu; the user may pick from a library of predesigned filter banks, may read in a filter bank of his own design, or may interactively design a filter bank with the software. There are choices for either wavelet or Modulated Lapped Transform (MLT) designs. In the case of a wavelet filter bank, the user specifies two parameters: the rank or number of channels M and the filter length L . The resulting orthonormal (paraunitary) wavelet filter has the maximal number of vanishing wavelet moments for the given filter length; these filter banks generalize Daubechies construction to the M -channel case, as described in section 4.4.1. For $M > 2$, the wavelet filters are determined by the characteristic Haar matrix (polyphase matrix evaluated at $z = 1$); this is set to the discrete cosine transform (DCT), because this choice has proven effective in applications. The other possibility for designing a paraunitary filter bank is the cosine-modulated or MLT case, in which one designs a prototype lowpass filter which is modulated via a DCT to uniformly partition the full frequency spectrum (section 4.5.3). Again, the user prescribes the rank M and the overall filter length

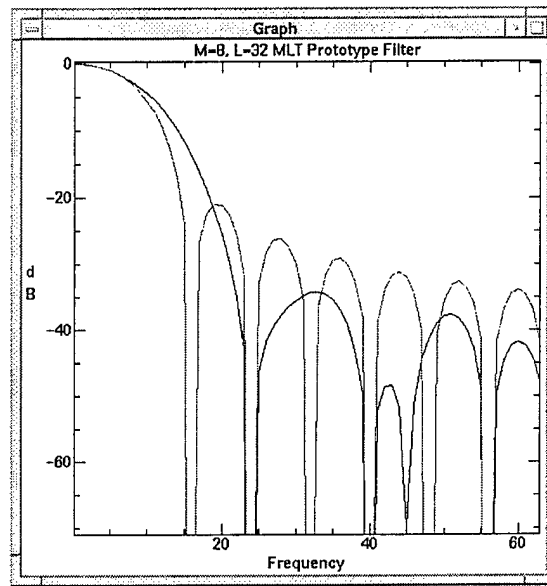


Figure 29: Two MLT prototype filter designs in WaveTool, with $M = 8$, $L = 32$.

L . When the filter length is greater than $2M$, there are free parameters which may be used to optimize the filter for superior stopband attenuation, narrow transition bandwidth, etc. Figure 29 shows two MLT prototype filter designs from the WaveTool for rank $M = 8$ and filter length $L = 32$ which trade off between height of the first sidelobe and transition bandwidth.

Once the user has designed a filter, he may then “accept” it into the nascent tree structure in the Tree Window. The filter bank coefficients (time domain impulse responses) are also saved to an ASCII file. The impulse and frequency responses of any filter bank in the tree may be displayed with a mouse click.

A salient feature of the WaveTool is the ability to create arbitrary tree structures of paraunitary filter banks. The classical wavelet/multiresolution decomposition [23], [64] is obtained by cascading rank 2 filter banks on their lowpass outputs only. Given any rank M filter bank in the Tree Window, the user may create such a Mallat tree with a single menu pick. However, non-Mallat trees have also proved quite useful in signal analysis (e.g. wavelet packet decompositions [22], the FBI’s WSQ fingerprint compression standard [7], and audio compression using a psychoacoustic model [98]). Any such tree may be created with the WaveTool.

Once the user has constructed a tree-structured filter bank, he or she can use it to operate on data. One loads in an input data file (in ASCII, binary, or MATLAB format); the input signal may be displayed in a graphics window. For example, we have loaded in the piecewise quadratic signal “nuquad” and displayed it, as shown in Figure 30. Plots such as this may be zoomed and relabeled at will, and dumped to PostScript output for inclusion in reports.

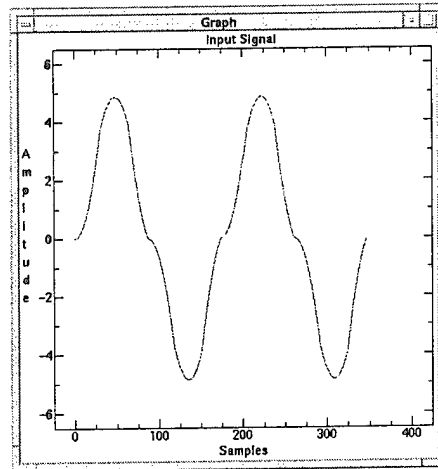


Figure 30: Piecewise quadratic input signal.

We then pick an “M2L6” wavelet filter from the Pick Filter menu (i.e. Daubechies’ rank 2, 6-tap orthonormal wavelet filter bank), and cascade it into a three-level Mallat tree, as shown in Figure 31. Notice that the Tree Window uses the convention that lowpass outputs are toward the top of the window. When we hit the Analyze button, we perform a decomposition of the input into the octave-based wavelet/subband decomposition. We can display the subband outputs as streaming from the output leaves of the tree-structured filter bank by hitting the Show Subbands button; this is also shown in Figure 31. The polynomial interpolation properties of the Daubechies 6-tap filter are shown clearly here – the only nonzero bandpass/highpass subband outputs occur at the knots of the piecewise quadratic input.

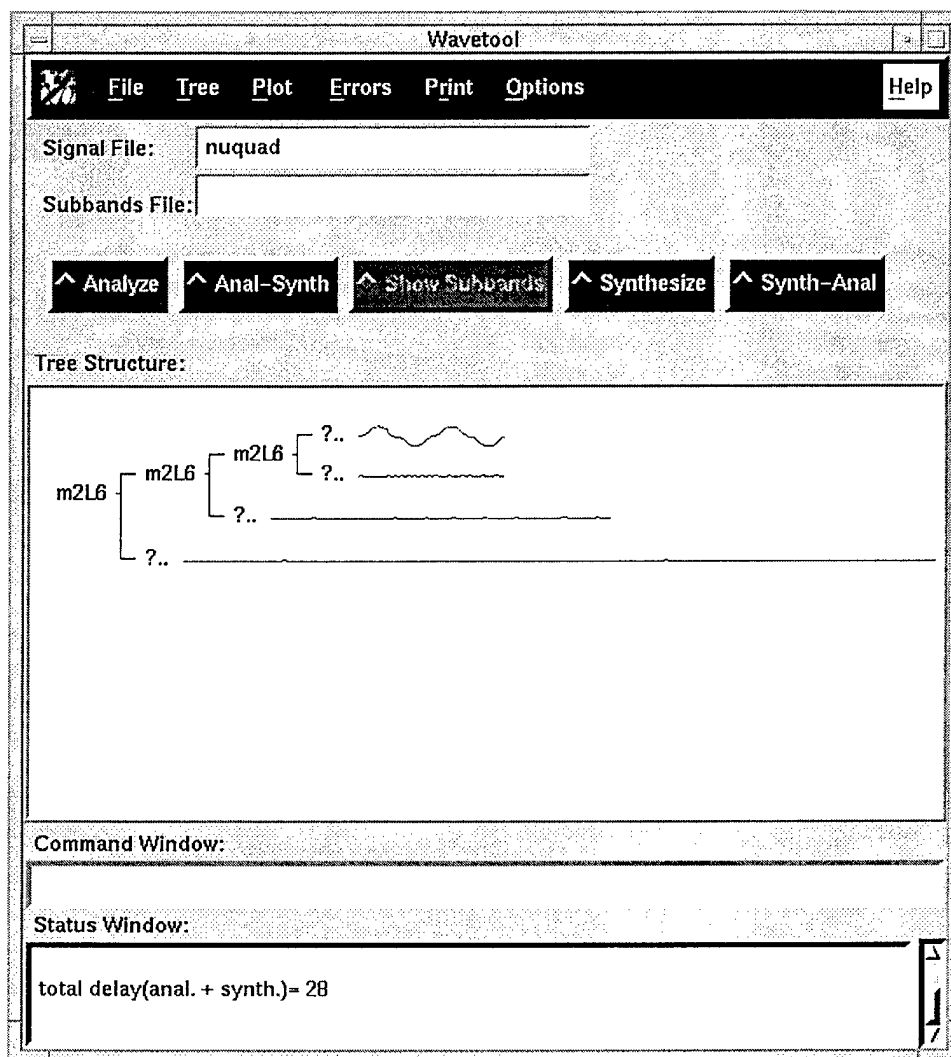


Figure 31: “D6” wavelet decomposition of the piecewise quadratic signal.

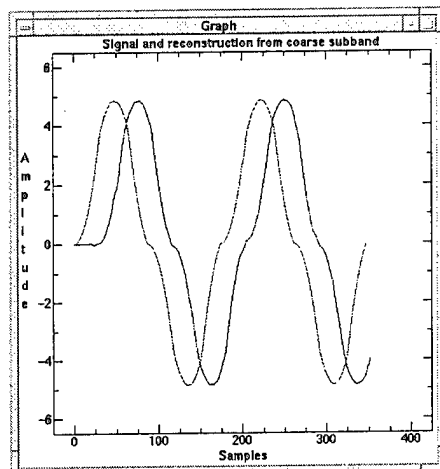


Figure 32: Piecewise quadratic input signal, and its coarse reconstruction.

Having applied the wavelet transform, one can then output the transform coefficients to ASCII or MATLAB data files. The MATLAB output option is particularly useful, for one can load the transform coefficients in as matrix variables and perform quantization, peak-picking, and other operations on the coefficients before saving back to file, and reloading into the WaveTool for the synthesis (reconstruction) operation. For example, let us take the wavelet subbands from our D6-Mallat tree decomposition of the piecewise quadratic signal, load them into MATLAB, and zero out the bandpass and highpass subbands. We then load the remaining coarse-scale transform coefficients back into the WaveTool and reconstruct into a signal. Figure 32 shows the original signal, followed by the reconstruction from the coarse-scale approximation at a delay of 28 samples. The coarse reconstruction hardly differs from the original piecewise quadratic input, again because of the interpolation properties of the Daubechies filters. Furthermore, one may measure signal reconstruction error in the WaveTool, using a variety of norms (ℓ^1 , ℓ^2 , ℓ^∞). For example, the coarse reconstruction of Figure 32 has an ℓ^∞ error of 0.1264093. These data input/output features are not the only means by which the WaveTool interacts with other software; once a wavelet/subband decomposition has been designed in the WaveTool, the algorithm may be incorporated into a larger MATLAB simulation (of, say, a compression or communications system) via MEX-file implementations of the filtering algorithms. Further details on the operation of the WaveTool may be found in the software's User's Guide [4] and Reference Manual [3].

In addition to its use in-house at Aware for tasks such as waveform design for Discrete Wavelet MultiTone modulation (section 6.1), the WaveTool software has been installed at a number of academic, government, and industrial beta sites, and ultimately released as a commercial software product. This is discussed in section 10.1 below.

6 Applications

We now turn our discussion to the third leg of this project – application of higher-rank wavelet techniques to telecommunications and signal processing. A surprise winner has been the use of rank M wavelet filter banks to multicarrier modulation, a leading technique for high-bitrate last-mile telecommunications, which we discuss first. We then examine the application of the wavelet transforms and design techniques to the compression of image data, including new data regimes such as seismic, multispectral, and sonar data. Finally, we review our application of emerging multiwavelet techniques to signal and image processing.

6.1 DWMT – Multicarrier Modulation via Rank M Wavelets

Multicarrier modulation has recently emerged as a superior method for achieving high-bitrate transmission over the “last mile” of the installed wire plant. This application is of particular importance, given the dramatic increase in use of the Internet and the attendant need for high-bandwidth connections to the home. The last mile usually takes the form of either twisted-pair copper wire or a hybrid fiber-coax cable. ANSI’s T1E1.4 committee has chosen a multicarrier scheme as the standard for Asymmetric Digital Subscriber Line (ADSL) transmission. ADSL is an emerging technology that promises upwards of 6 Mbits/sec into the home or business, using the existing copper plant, while simultaneously allowing conventional telephone service over the same line. Multicarrier can also be used for hybrid fiber-coax networks, to provide two-way communications services (data delivery or telephony) over the existing cable television (CATV) network.

Bitrates in the Megabit/sec range employ frequency bandwidths of 1 MHz or more; both twisted-pair and HFC have particular impairments across these wide frequency bands. In the case of twisted-pair copper, the available signal-to-noise ratio (SNR) for information throughput varies strongly as a function of frequency, particularly over long (10,000 feet or more) loops and those with “bridge taps” (unterminated, unused stubs of wire). Both twisted pair and HFC, but particularly the HFC upstream channel (5-40 MHz) are also susceptible to powerful narrowband interference, such as that caused by AM radio transmissions. As we discuss below, each of these impairments can be bypassed or overcome by a multicarrier modulation scheme.

In multicarrier modulation, the transmission channel is partitioned into a number M of subchannels (usually $128 \leq M \leq 512$), each with its own associated carrier [78]. This is usually accomplished digitally, using an orthogonal transformation. At the receiver, the inverse transform is performed to demodulate the data. This can be interpreted as transmultiplexer taking time-division-multiplexed (TDM) data into frequency-division-multiplexed (FDM) data. Thus multicarrier modulation provides an efficient means to access, transmit, and distribute multiple data streams. Each

subchannel has its own characteristic SNR that can be measured and is then used to determine the constellation size (number of bits supported) for that subchannel. In contrast to a single-carrier scheme, multicarrier gives a fine-grained (M grains, to be exact) decomposition of the transmission channel for the purpose of assigning bit levels. This allows for more nearly optimal use of the available bandwidth in the case of SNR that varies strongly with frequency, as in twisted-pair transmission. Multicarrier also provides superior immunity to impulse noise (compared to single-carrier systems), and is also particularly effective at combatting narrowband interference. A subchannel that is affected by a narrowband interferer can be easily identified, and the data rate in that subchannel reduced.

Early multicarrier schemes [124], [54] used a Fourier transform (DFT) as the orthogonal modulating transform. This has the advantage of fast computational algorithms, but also the weakness of large spectral overlap among the frequency responses $A_k(\omega)$ of the distinct signalling waveforms. This can lead to substantial intersymbol interference, and a lack of robustness to narrowband interference.

Aware, Inc. has pioneered the use of wavelets and multirate filter banks for multicarrier modulation [111], [90] [88]. This approach, called Discrete Wavelet MultiTone (DWMT) or overlapped multitone modulation, yields superior bandwidth utilization and robustness in the face of impulsive and narrowband noise [127]. The block Fourier transform is replaced by an orthogonal overlapped transform (a rank M wavelet matrix with genus $g > 1$). The wavelet orthogonality conditions (2) prove to be equivalent to zero intersymbol and inter-channel interference (ISI and ICI) among the signalling waveforms. The use of an overlapped transform enables a tradeoff between time duration of the transform and spectral isolation of the subchannels; even a low degree of overlap yields significant improvements over the block (nonoverlapping) Fourier case. The superior subchannelization offered by the filter bank leads to dramatically superior performance in the presence of narrowband (e.g. radio-frequency) interference.

6.1.1 Optimization of wavelet algorithms for DWMT

The various techniques discussed in section 4 for wavelet filter design and computation have a direct application to DWMT modulation. In particular, we have found the cosine-modulated filter banks (Section 4.5.3) most useful as an overlapped orthogonal transform for multitone modulation. Cosine-modulated filter banks are determined by a single lowpass “prototype” filter. This simplifies the filter bank design problem, enabling rapid optimization of the wavelet filter bank for properties such as subchannel isolation. The notion of genus or overlap provides a handy parameter for the tradeoff between filter bank latency (and computational complexity) and narrowband interference rejection, for a given number of subchannels M . This is illustrated in Figure 33, which compares two different wavelet filter bank responses (sets of tones)

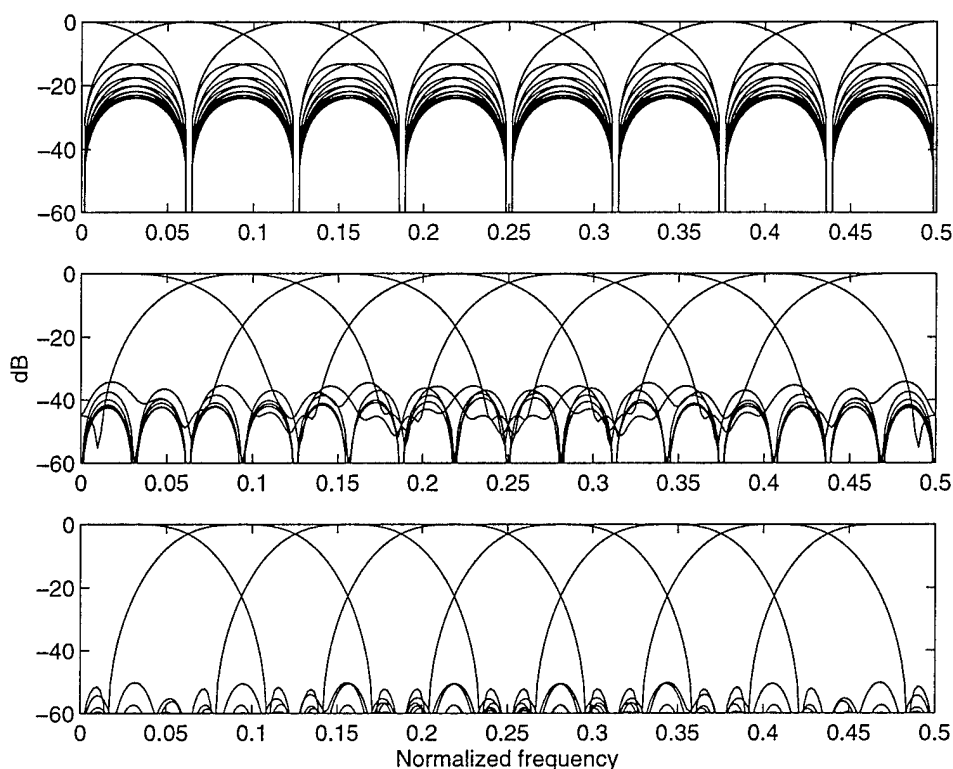


Figure 33: Subchannel frequency responses for a multitone system with $M = 8$. Top: nonoverlapped system based on the DFT. Middle: overlapped multitone system with genus $g = 4$. Bottom: overlapped multitone system with genus $g = 8$.

with those of a Fourier transform.

The top plot in Figure 33 shows the frequency responses of the DFT (used in conventional discrete multitone modulation) with $M = 8$. The peak sidelobe is at -13 dB, with $1/f$ decay. The second plot shows the subchannel frequency responses of an $M = 8$ genus $g = 4$ cosine-modulated wavelet filter bank, as used in DWMT. This transform has sidelobes below -35 dB. The third plot, of an $M = 8$ genus $g = 8$ transform shows sidelobes below -50 dB, clearly demonstrating the superior subchannelization offered by the overlapping wavelet transform. The wavelet filter banks displayed in this example were designed using the QCLS algorithm (section 4.4.7) as implemented in the WaveTool software.

The structures of section 4 also lead directly to fast algorithms for wavelet-based modulation. In particular, the cosine-modulated filter banks can be realized as a combination of a Discrete Cosine Transform (DCT) and a “polyphase windowing” step. Making use of previous work on fast algorithms for the DCT, the entire cosine-modulated filter bank computation may be done at a cost not much greater than that of the DFT with the same number of subchannels. Details appear in [110], [65], [66].

6.1.2 Hardware demonstration of a DWMT modem

We have developed a prototype DWMT modem and made laboratory measurements that demonstrate the superior performance of the system in the presence of narrow-band interference. Figure 34 shows a block diagram of the prototype system. The input data consists of a serial TDM data stream that is divided into frames of length B bits. Given an input data rate of R_b and a multicarrier frame duration T , the number of bits per frame is $B = R_b T$. For the prototype system, the parameters are $R_b = 2.048$ Mbits/sec, $T = 125 \mu\text{sec}$, and $B = 256$ bits/frame. Nominal data rates on the order of 5 Mbits/sec are achievable and a maximum rate of approximately 8 Mbits/sec can be obtained in a 2 MHz band.

The data bits are encoded into multilevel PAM symbols and the PAM symbols are then orthogonally mapped to individual frequency subchannels via the inverse wavelet transform modulator. Note that the subchannels are grouped into pairs and modulated with the equivalent of a QAM constellation. The wavelet transform is based on a single rank M wavelet matrix (corresponding to M subchannels) of genus g . The properties of the transform are such that the subchannels overlap spectrally and the pulses transmitted in a subchannel overlap in time, while orthogonality among all symbols is maintained. The frequency overlap of adjacent subchannels results in spectrally efficient transmission and the time overlap of the pulses provides spectral shaping of the individual subchannel filters. The transform modulator produces a time domain sequence that is output to the D/A converter. In the downstream direction, the signal is broadcast to all modems and in the upstream direction, the analog signals from the subscriber modems are power combined and then transmitted to the head end. The experimental setup consisted of an head end transceiver (as would be found at an Optical Network Unit) and 3 remote transceivers (Figure 35).

At the receiver, the baseband analog signal is filtered and digitized. The digital time domain signal is transformed back into the PAM symbols via the wavelet transform. The demodulating transform is the "analysis" filter bank based on the same rank M genus g wavelet matrix as the modulating transform. At the receiver, this provides a set of matched filters with respect to the modulating transform. An equalizer is used to correct for channel distortion and timing errors. After equalization, the PAM symbols are decoded into bits and the original data stream is reproduced by the parallel-to-serial multiplexer. The receiver can be constructed to demodulate all of the subchannels to recover the complete data stream or a specific subset of the subchannels for information addressed to individual users. Dynamic allocation of subchannels and transmission capacity can be implemented based on user demand.

The measured results are presented in Figures 36 through 41. In Figure 36, the spectrum of the transmitted signal with 64 QAM modulation on the payload tones is given. In addition to the group of tones used for payload (P), a set of tones at the upper and lower band edges are designated for timing and ranging (TR) along with

a group for control signaling (C). The data rate is 2.048 Mbits/sec and the entire payload signal is confined to a bandwidth of approximately 425 kHz (see markers on the spectral plot). The bit error rate (BER) was measured as a function of the signal-to-noise ratio (SNR) for both 16-QAM and 64-QAM constellations (Figure 37). The measured results closely track the calculated (solid line) values. The calculations are based on standard QAM performance modeling in the presence of Gaussian noise [78] and the SNR is measured digitally at the decision statistic after equalization.

Figure 38 shows the spectrum of the signal (which appears white over the selected frequency band) and inserted radio frequency interference (RFI). The SNR for each tone was measured over the band for different power levels of RFI ($P_{ingress} = -47, -35, -22$ dBm) with the signal power per tone held constant ($P_{tone} = -41$ dBm). Figure 39 graphs the measured results, demonstrating how the DWMT waveform provides excellent isolation of the ingress noise. Even under the worst case RFI level, only 5 tones (approximately 20 kHz of the spectrum) were degraded by the noise. The remaining tones outside of this band attain the same SNR whether or not the RFI is present.

After injecting the RFI directly in-band as shown in Figure 38, a group of tones were reallocated to a clean portion of the channel and measurements were made as a function of the frequency difference from the center of the closest DWMT tone to the center of the RFI. The spectrum of the corresponding signal with 64 QAM modulation levels and ingress noise (RFI) spectra are depicted in Figure 40. With no RFI, the SNR for each DWMT tone is about 35 dB. Note that the BER was measured with the system subject to the noise levels in Figure 40 (signal/ingress frequency difference is 14 kHz). No errors were recorded over a 24 hour period, indicating a BER less than 1×10^{-11} . As the ingress noise was moved closer to the signal, the BER degrades as shown in Figure 41. The measured results again clearly demonstrate the ability of the DWMT system to isolate the ingress noise. Changing the ingress frequency by 2 kHz reduces the BER from 10^{-9} to 10^{-3} . As the power level of the ingress noise is increased, the corresponding frequency separation between the DWMT signal and the ingress must also increase. This is a direct result of the ingress spectrum that consists of a main lobe and a series of side lobes.

In conclusion, wavelet filter banks have been shown both theoretically and experimentally to add a new dimension to multicarrier modulation. The use of overlap (as measured by the genus parameter g) provides a means for tradeoff between subchannel isolation and system latency and computational requirements. The subchannel isolation provided by wavelet modulation yields increased robustness and flexibility with regard to narrowband noise. Fast algorithms for wavelet computation map directly to the multicarrier modulation application, minimizing the additional cost due to the use of an overlapped transform. DWMT is a very promising technique for high-bandwidth transmission over both twisted-pair copper wire and hybrid fiber-coax wirelines in the "last mile" of network connections to home, school, and business.

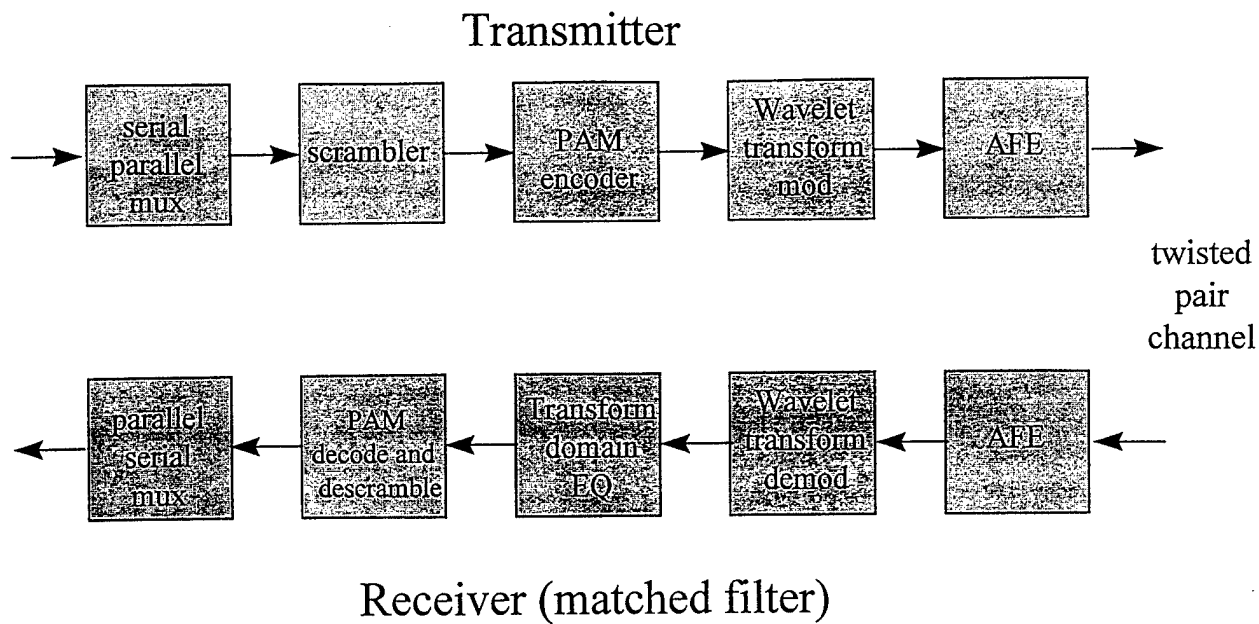


Figure 34: Schematic diagram of a Discrete Wavelet Multitone (DWMT) system.

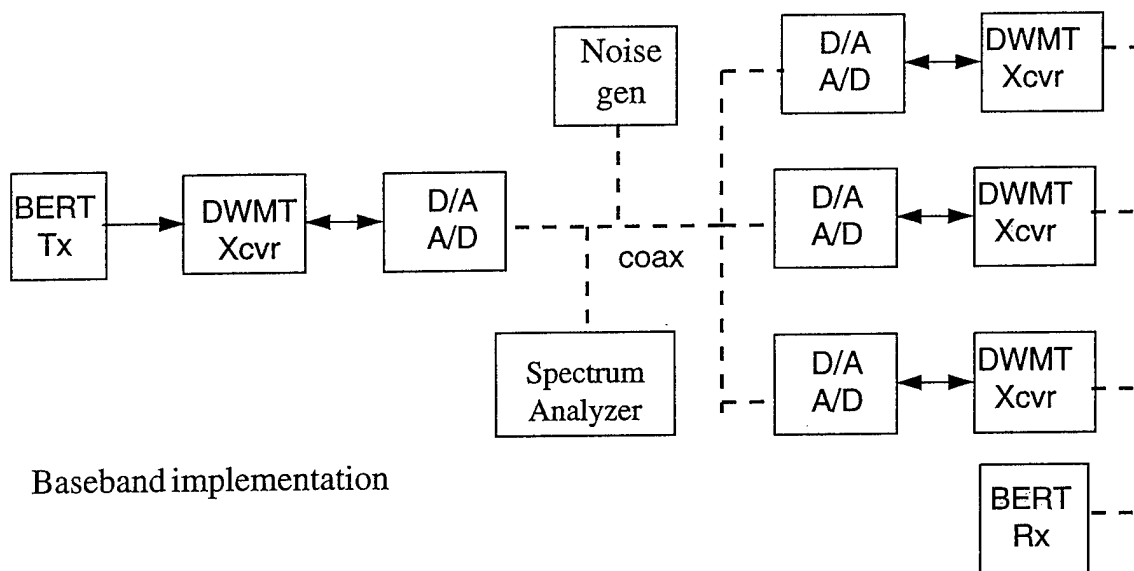


Figure 35: Experimental baseband DWMT setup with 1 head end modem and 3 remote modems.

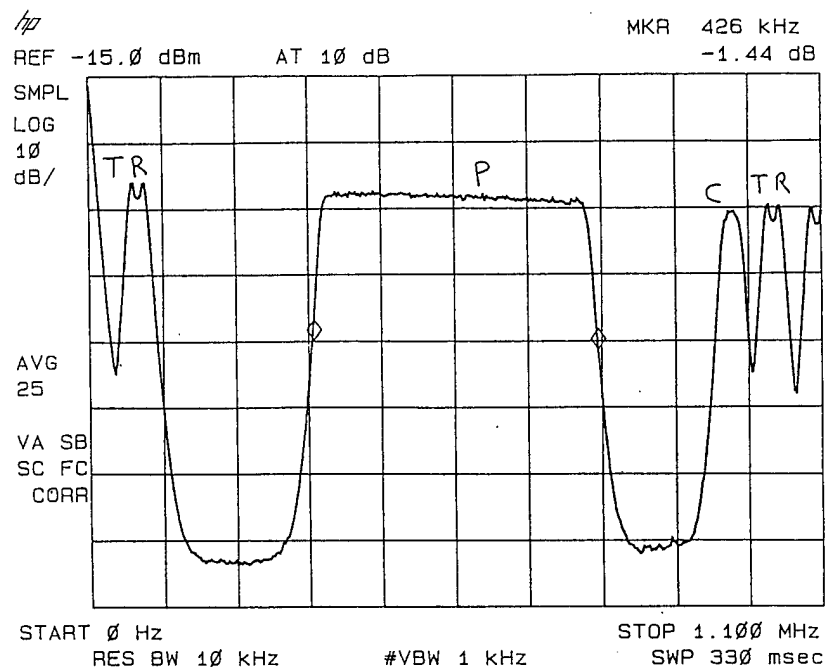


Figure 36: Measured spectrum of a downstream DWMT signal using 64 QAM modulation for the payload tones. In addition to the tones designated for payload (P), additional tones are used for ranging (R) and control (C).

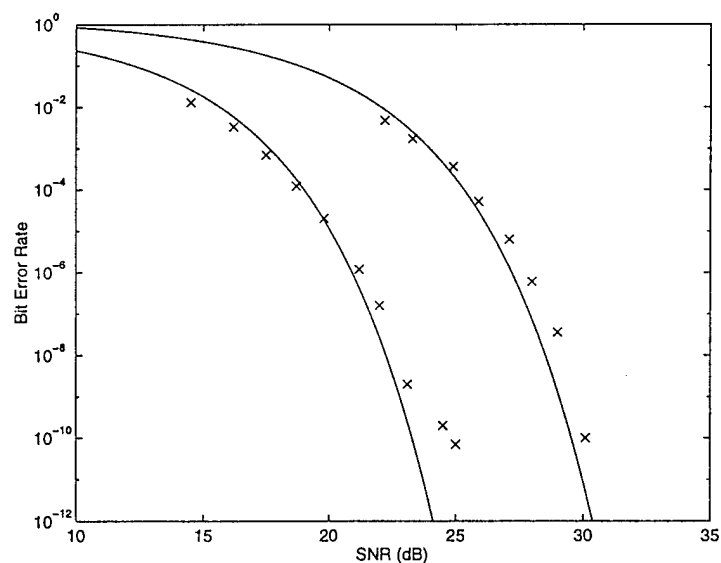


Figure 37: Measured bit error rate (BER) vs. SNR for 16-QAM and 64-QAM transmission with DWMT. Solid lines are the theoretical predictions.

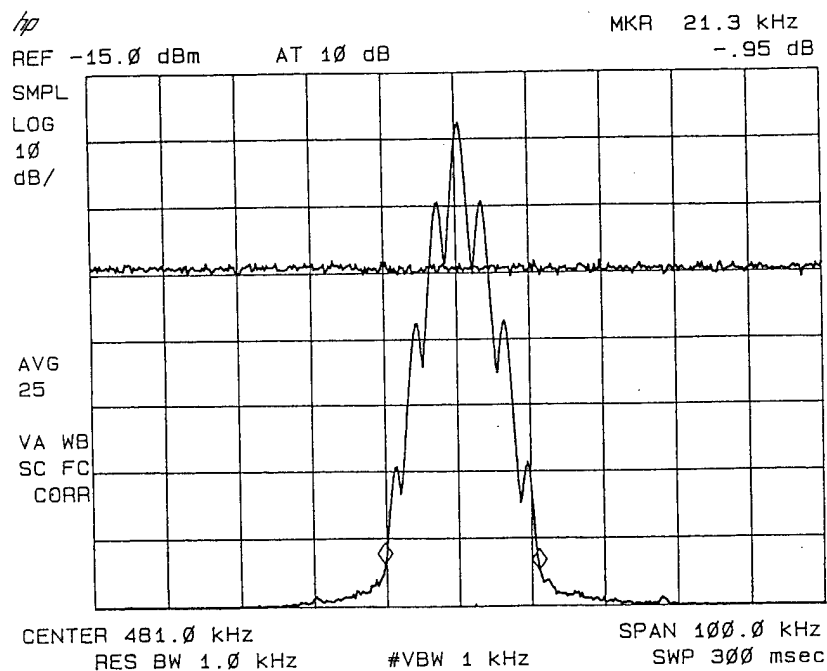


Figure 38: Spectral plot of a DWMT signal (appears as a white spectrum) with RFI present.

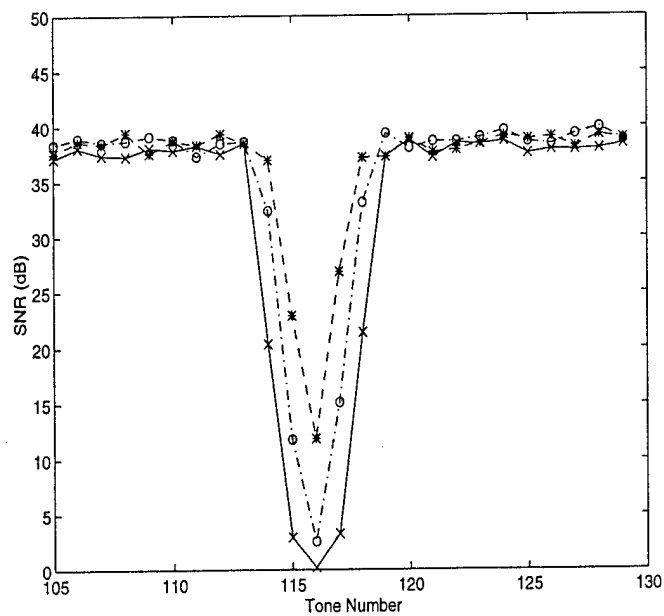


Figure 39: Measured SNR for each tone with RFI at power levels of -47 dBm, -35 dBm, and -22dBm. Even with very high RFI levels, only 5 tones are degraded by the noise.

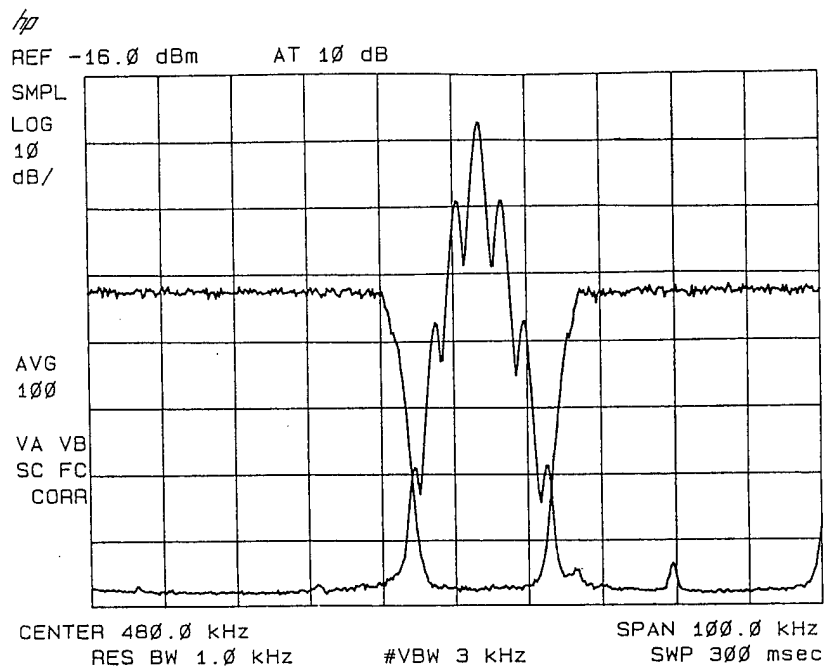


Figure 40: Spectral plot of a DWMT signal with tones reallocated to avoid the RFI.

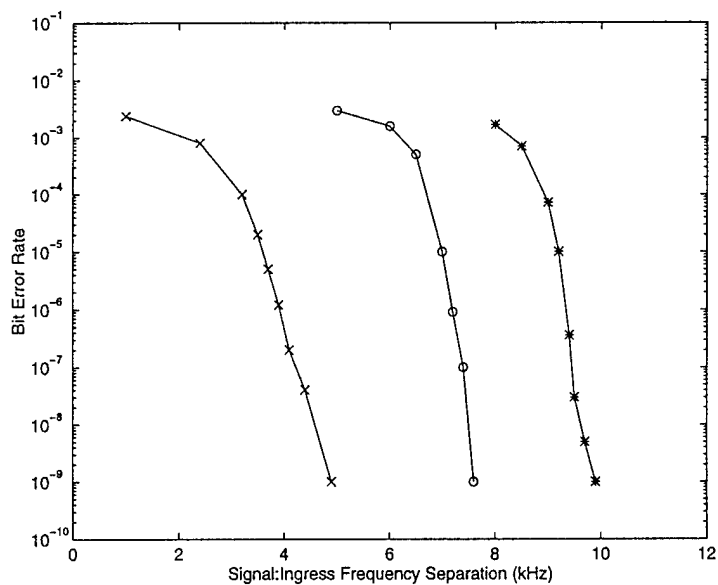


Figure 41: Measured BER for the noise conditions shown in Figure 40 for three different levels of RFI and 64 QAM signalling. As center frequency of the RFI is varied to bring the RFI within the band of a tone, the BER degrades.

6.2 Image Compression

Image data compression has been one of the first and foremost applications of wavelet methods. We continued Aware's efforts in wavelet-based image compression as part of the higher rank wavelet project. Our work involved applying the new rank M wavelet transforms to image compression, fine-tuning the quantization algorithms, and applying wavelet techniques to the compression of new types of data, including seismic data and multispectral LANDSAT data.

6.2.1 Review of wavelet-based image compression

We begin with a review of wavelet-based image compression. A typical lossy compression algorithm consists of three basic steps: a reversible lossless transform, a lossy quantizer and a standard lossless encoder. The transform step may consist of a Discrete Cosine Transform (DCT) as in the ANSI standard JPEG compression algorithm [75], or a wavelet transform as in Aware's AccuPress production software, and in the research reported here. The energy compaction properties of the transform domain representation is crucial to the overall success of a compression algorithm. While the DCT provides a compact data representation composed of stationary harmonics, wavelet transforms provide a superior representation for data composed of sharp edges and other transient characteristics as well as broad smooth areas. The quantization step, which accounts for all of the "loss" in the compression, seeks to represent the more significant components of the transform representation with proportionally more accuracy or bits than the less significant components. The ability of the wavelet transform to compactly segregate the important parts of seismic signals allows the quantizer, for a given bitrate, to cause significantly less "loss" than would be the case for other representations such as a Fourier transform [64]. Because the quantizer output is put through a lossless entropy coder to create the final compressed bitstream, we employ an entropy-constrained quantization technique. The quantized data is then passed through a combined Huffman and zero-run-length encoding to approach the true entropy of the quantized bitstream.

6.2.2 New wavelet transforms

A significant piece of the work carried out in this project was the development of new families of wavelet transforms (the rank M wavelets). We applied several of the new wavelet constructions to image compression using Aware's standard wavelet compression algorithm; results were briefly quoted in sections 4.4.4 and 4.5.2. Further details can be found in references [45], [48], and [44].

6.2.3 Improved quantization techniques

Our efforts on improving (lossy) quantization were twofold: an examination of so-called “dead zones” for wavelet scalar quantizers, and the development of an operational rate-distortion-optimal uniform quantization scheme.

Empirical evidence [33] has suggested that the subjective and objective performance of a uniform quantizer may be improved in transform coding applications by allowing the center decision region in a uniform quantizer (also known as the *zero bin*) to be wider than the other decision regions by a constant factor. In our notation, the zero bin width is given by

$$\Delta_0 = \beta \Delta$$

where Δ is the uniform quantizer bin width.

The expanded zero quantization bin is sometimes called the *dead zone*. It is thought that the perceptual improvement afforded by an expanded zero bin is due to the fact that the contribution of basis functions with small coefficients is underemphasized rather than exaggerated. While this does not minimize the mean squared error, it improves perceptual quality because small coefficients increased by amplitude quantization manifest themselves perceptually as noise. With an expanded dead zone quantizer, more of the extremely small coefficients are quantized to zero, yielding a more visually pleasing reconstruction and aiding subsequent zero-run-length coding.

We conducted experiments with varying dead zone widths on a variety of image data, ranging from natural images to overhead satellite images to seismic datasets. We found that using an expanded center bin yielded an increase in peak signal-to-noise ratio (PSNR) between 0.2 to 0.5 dB. For natural and satellite imagery, the maximally effective dead zone width (around $\beta = 1.8$) yielded a PSNR improvement of about 0.3 dB. Both the location of the peak and the amount of improvement seem to be typical for natural images such as the NITF series at this compression ratio. As the compression ratio increase, the peak PSNR improvement occurred for larger values of β . The most significant difference between the seismic data and the natural images was that the peak SNR improvement seemed to occur at a much lower β for the seismic data. In order to compress both types of image data effectively, we chose a fixed β of 1.4, which gave reasonable performance improvement for both data types. Further details appear in [17].

In addition to this examination of dead-zone quantizers, we developed an operational rate-distortion-optimized quantizer based on the methods of Shoham and Gersho [93], applied to the wavelet transform domain. It has been observed [108], [64] that histograms of the subband outputs of the wavelet transform applied to image data obey a generalized Gaussian distribution. Beginning with this assumption, we employed an entropy-constrained uniform quantizer with a mid-rise transfer curve.

The bit-allocation algorithms employed [93], [80] require knowledge of the rate-distortion characteristics of the quantized data. Since our quantizer is followed by an

entropy encoder, we make the assumption that the source will be compressed down to its entropy level. For this reason, we examine the entropy of various quantized sources versus the distortion rather than rate versus distortion. We also chose uniform quantizers as opposed to the pdf-optimized Lloyd-Max quantizers ([108], [55]).

First, we compute the entropy of the quantized source as a function of the bin width Δ and the generalized Gaussian exponent α . The entropy is defined to be

$$H = - \sum_{i=-L}^L p_i \log_2 p_i$$

where

$$p_i = \int_{d_i}^{d_{i+1}} f_X(x; \alpha) dx \quad i \neq 0$$

and

$$p_0 = 2 \int_0^{d_1} f_X(x; \alpha) dx$$

that is, a sum of integrals of the probability distribution of the source over the i^{th} decision region. Using these expressions, we can easily compute the value of the entropy as a function of the stepsize Δ and the subband variance σ .

The *distortion* is defined to be the mean squared error introduced by the quantization process. It is clear that as the step size becomes finer and finer, the approximation is increasingly better, thus yielding smaller distortion. On the other hand, the more finely quantized the data is, the more bits are needed to represent it. This is the fundamental tradeoff involved in the lossy transmission of information — the rate/distortion tradeoff.

The distortion due to quantization in a given subband is

$$D = E [(X - Q(X))^2] = \int_{-\infty}^{\infty} (x - Q(x))^2 f_X(x; \alpha) dx$$

Since the quantization bins form a partition of the real line, we may write the integral as a summation of integrals over each quantization bin separately plus the integral over the tails of the distribution.

$$D = \sum_{i=-\infty}^{\infty} \int_{d_i}^{d_{i+1}} (x - r_i)^2 f_X(x; \alpha) dx$$

where r_i is the i -th reconstruction level. Using the component parts of this equation, we can compute the overall distortion introduced by the quantizer. Having now parameterized both the entropy H and the distortion D by the binwidth, subband variance, and Gaussian exponent, we can now compute parameterized tables of entropy versus distortion for use in quantization and bit allocation.

Given such precomputed rate-distortion tables for a fixed set of generalized Gaussian exponents, at runtime we measure the subband variance and kurtosis to determine the best-fit Gaussian exponent to each subband. This determines which

rate-distortion tables are used to determine the optimal bit allocation among the subbands, using the operational rate-distortion methodology of [93], which follows the convex hull of the rate-distortion curve.

The resulting quantizer yields significantly superior performance for wavelet compression, when compared with simple scalar quantizers such as that of the FBI's Wavelet Scalar Quantization specification [7]. Further details are given in [18].

6.2.4 Seismic data compression

In addition to improvement of our basic wavelet transform compression algorithm, we applied wavelet compression to several new types of data. One of the most promising developments has been the application to seismic data compression. In the seismic regime we tested a variety of wavelet bases and tree structures on a variety of seismic data types. After a brief review of the specific areas of experimentation, we explore our prospects and successes in commercialization of wavelet-based seismic data compression.

One of our first areas of experimentation was to try different rank 2 basis functions (filters) and different multiresolution trees on sample seismic datasets. The Daubechies biorthogonal 7/9-tap filters [11] were compared to Daubechies 6,8,10, and 12 tap orthogonal filters. The biorthogonal 7-9 tap filters provided up to a several dB increase in pSNR (peak signal-to-noise ratio) over any of the Daubechies orthogonal filters at compression ratios of between 20:1 and 80:1. This is consistent with the results of Macq and Mertens on natural images [63].

We also explored the use of new families of rank M wavelets; in particular, the rank M linear-phase orthogonal wavelets discussed in section 4.5.2. Compression results (based on the size of the entropy-coded bitstream) are shown in Table 3. We compared a 5-level Mallat tree based on the rank 2 Daubechies (7,9)-tap biorthogonal filter pair [11] with a 3-level Mallat tree [64] based on the rank 4 12-tap GenLOT shown in Figure 25. The rank 4 wavelet offered superior or comparable performance across all compression ratios; peak SNR's were approximately 1 dB greater for the rank 4 linear-phase wavelet.

	8:1		16:1		32:1	
	pSNR	Max	pSNR	Max	pSNR	Max
$M = 2$	51.6	1101	40.7	4262	32.7	12544
$M = 4$	52.6	1025	41.3	4298	33.7	12260

Table 3: Peak SNR and maximum errors for compression of 2-d seismic data example (Mallat tree, $M = 2$ (7,9)-tap pair and $M = 4$ GenLOT).

We also tested a number of non-Mallat transform trees to evaluate the merit of us-

ing hardcoded trees for a seismic-specific wavelet compression algorithm. The recently released Federal Bureau of Investigation specification for fingerprint compression has such a hardcoded tree which was found to be optimal for fingerprints (Figure 27). We evaluated several non-Mallat trees on four very different types of seismic data. Our initial conclusion is that the wide variation in seismic data types makes the use of a single fixed non-Mallat tree impractical. However, a fast adaptive tree could be successful for seismic applications.

The efficiency of a standard two-dimensional wavelet transform for data which is moderately rectangular was found to be generally adequate. Seismic data tends to be grouped into logical frames with approximate dimensions of 240 columns by 2000 rows. We compared a standard Mallat tree with an initial rowwise transform followed by a Mallat tree on the low pass output and found the results to be comparable.

Seismic data can be grouped into two broad categories, stacked and pre-stack. Stacked data contains a large percentage of horizontally aligned energy caused by the generally flat structure of the earth's sedimentary rock layers. Pre-stack data is composed largely of hyperbolically aligned energy due to the geometry of the seismic data acquisition process. In theory the horizontally aligned arrivals should be easier to compress, but in practice the hyperbolically aligned seismic arrivals were preserved as well as the horizontally aligned arrivals after passing through the same wavelet-transform based compression algorithm. This was the opinion of the seismic experts who examined the data.

Another area of algorithm development lay in the use of companding. Conditioning or companding the histogram of data prior to compression is a known method for improving compression performance [55]. This is due to improved quantizer performance because the companding operation has in some way caused the quantizer to preferentially preserve amplitudes which are important to the application at hand. The largest amplitudes in a seismic data set are usually the least important components of the data and therefore should not be preserved at the expense of other more important components. While it is not possible to simply compress the reciprocal of the seismic data, it does help to condition the data prior to compression to help boost small but important features. We experimented with standard logarithm based companding curves and found moderate improvement at low compression ratios but unacceptable results for compression ratio greater than 15:1. Much better results have been achieved by allowing the seismic data processing specialist to do this conditioning based on some physical model of the earth prior to compression. Specific characteristics of the exploration regime can be used in the conditioning of the data, rather than a simplistic blind companding curve. Initial results indicate that conditioning of seismic data is an important step prior to compression, particularly for pre-stack data.

We see considerable commercial potential for wavelet-based seismic data compression. A typical seismic acquisition boat collects between one and ten gigabytes

of seismic data per day. Historically, the seismic industry has been the initial consumer very large mass storage and data transmission products and continues to do so. Aware has worked with a number of companies involved in actual oil and gas exploration as well as in data acquisition and data processing. Lossy data compression is a novel concept to the entire seismic industry and has been quite well received. We have had experts in seismic data processing state that the compression results at 10:1 are virtually "lossless" and 20:1 is sufficient for the vast majority of processing. We have successfully integrated compression into data transfer systems which use the Inmarsat communications satellites. These systems allow technicians on seismic boats to send samples of seismic data via Inmarsat back to a home office where experts can examine the data and determine whether changes need to be made to the boat's acquisition plans. While the file transfer systems previously existed, it was virtually impossible to send seismic data due to the cost of the satellite link. Data compression has had an enormous impact on costs, enabling the practical economic transmission of seismic data from acquisition vessels in near real time. We are also working on using compression to reduce seismic mass storage costs in on-shore processing centers. Our close contact with the seismic industry has allowed the algorithmic efforts detailed above to be reviewed by a demanding and knowledgeable audience. Further details appear in the publications [14], [83].

6.2.5 Compression of Multispectral (LANDSAT) Data

As part of the contract work effort, we also extended lossy wavelet image compression methods to 12-band multispectral LANDSAT satellite imagery. Multispectral and hyperspectral images, usually gathered by satellite, consist of many two-dimensional images. Each two-dimensional image corresponds to a narrow spectral band. The LANDSAT data we worked with had 12 such spectral planes or bands; next-generation satellites capture hyperspectral images with 60 to 100 planes. Each band is usually about 1000×1000 pixels, with 8 bits per pixel. There is significant correlation among the spectral bands; exploiting this correlation is key to successful multispectral data compression.

In practice, the multispectral data will be gathered and compressed on board a satellite, then transmitted to a ground station for storage, decompression, and retrieval. The satellite will have strict limitations on power consumption and hence the possible complexity/memory requirements of the compression algorithm, and must operate in real time. In contrast, the ground station will have practically nonexistent power limitations, and the received or stored data will often be browsed in a non-real-time setting. In view of this asymmetrical system, several requirements appear. The compression algorithm must not have too high a complexity. It must also be of high quality, for the satellite cannot go back and recapture data that has been found to be of low fidelity. Furthermore, a scheme which lends itself to searching and browsing is

preferable.

A wavelet approach wins on several of these counts. Wavelet compression has consistently proven itself to be of higher quality than comparable DCT-based algorithms [16], while of only slightly greater complexity. Because the wavelet transform is a multiresolution representation, it offers several advantages for LANDSAT applications. A few coarse-scale coefficients can efficiently describe broad subregions of a large satellite image while yielding high compression ratios. This is in sharp distinction to the blocking artifacts and limited effective compression ratios of block-DCT algorithms such as the JPEG standard [75]. Wavelet data structures also naturally lend themselves to multiresolution image browsing, allowing the user to call up a coarse-scale representation of an image and only request finer detail when he actually needs it.

Our experiments involved the exploration of several alternatives in the design of a compression algorithm for multispectral data, within the constraint of employing a wavelet transform in each of the two-dimensional bands of the multispectral dataset. The LANDSAT test imagery was of varying sizes, roughly 900×1100 pixels of 12 bands by 8 bits. We applied a wavelet transform in the spatial dimension, consisting of a six level Mallat tree of Daubechies' rank 2 linear-phase (7,9)-tap biorthogonal wavelets [11], yielding a nonexpansive transform of the data. We then considered the following alternatives:

- Independent wavelet compression processing of each spectral band.
- Independent wavelet transforms of each band, followed by global bit allocation (quantizing all subbands of all spectral bands at once).
- The use of a Discrete Cosine Transform (DCT) in the 12-band spectral dimension as well as wavelet transforms spatially, followed by global bit allocation.
- The use of the Karhunen-Loeve Transform (KLT) in the 12-band spectral dimension as well as wavelet transforms spatially, followed by global bit allocation.

The first two alternatives were considered as a baseline, against which to measure the effectiveness of the spectral DCT and KLT transforms in exploiting cross-band redundancy. The only difference between the first two approaches is in the bit allocation, where the first employed distinct bit allocations for each band, while the second approach allocated bits across all the subbands. Not surprisingly, the two approaches displayed very similar performance.

We then considered the latter pair of alternatives, performing either a Discrete Cosine Transform or a Karhunen-Loeve Transform across the 12 spectral bands. The KLT, which is the transform that diagonalizes the cross-correlation matrix of the data, is precisely that block transform that yields optimal energy compaction of the data. Because of this data-dependent optimality, the KLT will do the best possible

job of transforming the cross-band data for subsequent quantization and compression. However, we considered the DCT as well because system requirements may dictate the use of a fixed (non-data-dependent) transform with fast algorithms for on-board compression processing. The DCT is widely used, because it both possesses fast computational algorithms and serves as an approximation to the KLT for AR(1) processes [82].

Rate-distortion curves for the various approaches across a range of bitrates from .25 bits/pixel to 2 bits/pixel (compression ratios ranging from 32:1 to 4:1) are shown in Figure 42, while compressed/decompressed images are shown as Figure 43. Observe the significant decrease in distortion when one moves from compression based on a spatial-only transform to the addition of a cross-band transform such as the DCT or KLT. Interestingly, the DCT yielded 75% of the performance gain achieved by going from no spectral transform to the full-blown KLT.

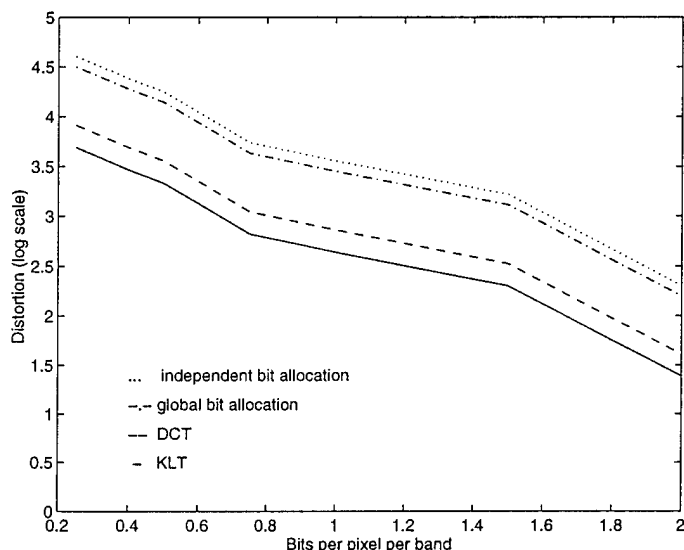


Figure 42: Rate-distortion curves for the four multispectral compression algorithms.

In conclusion, these experiments (reported in [41]) confirm the utility of wavelet transforms for multispectral image compression, and the importance of exploiting cross-band correlation. We also validated the use of a DCT as a stand-in for the KLT in exploiting cross-band correlations in a computationally efficient manner. Areas for future work include trying simple (e.g. Haar) wavelet transforms across the bands, employing zerotree structures [92] in the quantization and coding, and trying out ideas from video compression like motion estimation and compensation for the cross-band processing.

Figure 43: See attached page.

6.3 Sonar Data Compression in Real Time

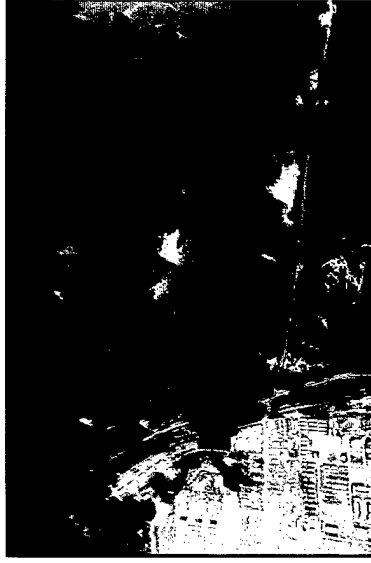
During 1993 Aware took part in an AntiSubmarine Warfare simulation/exercise using the Internet that was sponsored by ARPA's Maritime Systems and Technology Office [25]. Our role was to provide high performance data compression to reduce inter-site communications bandwidth requirements in a distributed simulation environment. Aware has already developed robust CD-quality real-time audio compression software [98, 99] which is based on rank M wavelets and multirate filtering. A critical part of sonar compression project was to tailor the basis choice to sonar signals, rather than CD audio. The WaveTool software proved invaluable in this task.

Given sample sonar data, the aim was to find a wavelet/subband decomposition which provided maximum energy compaction, i.e. localized the signal energy in as few subbands as possible. We used a well-known bit allocation formula that optimizes coding gain ([116])

$$b_k = b + \frac{1}{2} \log_2 \frac{\sigma_{x_k}^2}{\left(\prod_{i=0}^{M-1} \sigma_{x_i}^2\right)^{1/M}} \quad (43)$$

to achieve an overall bitrate b . We decided to restrict ourselves to the use of cosine-modulated perfect reconstruction filter banks, because of the available DCT-based fast algorithms for their computation, as well as the superiority of off-the-shelf filter design methods. We ran several sonar test signals through the Wavetool, testing filter banks of rank 4, 8, and 16. In each case, we analyzed the input signal, wrote the subband data out to file, quantized using the scheme (43), and then read the quantized data back in to the Wavetool and synthesized a reconstructed signal. Using the error measurement features of the Wavetool, we were able to evaluate mean-squared error and maximum error, and found the rank 8 filter banks to be optimal for the signals in question. We then varied the overlap of the $M = 8$ filter bank and decided on an overlap $N = 4$ design. Longer overlaps led to diminishing returns at a nonzero computational cost. The sonar data to be compressed was classified as active and passive. The passive data was principally lowpass, and using the Wavetool, we found that we could attain further compaction of the passive signals by iterating on the lowest channel of the $M = 8$ split with a rank 2 split. For this filter we used a 64-tap high stopband attenuation perfect-reconstruction bank. The facilities provided by the WaveTool led to a significant reduction in algorithm design time for the sonar data compression demonstration.

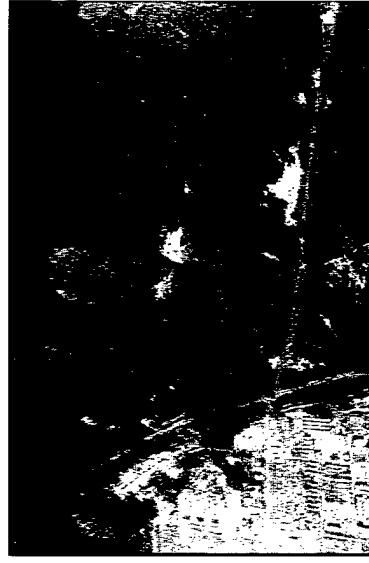
Figure 43: LANDAT image, original and compressed to .25 bits/pixel (32:1).



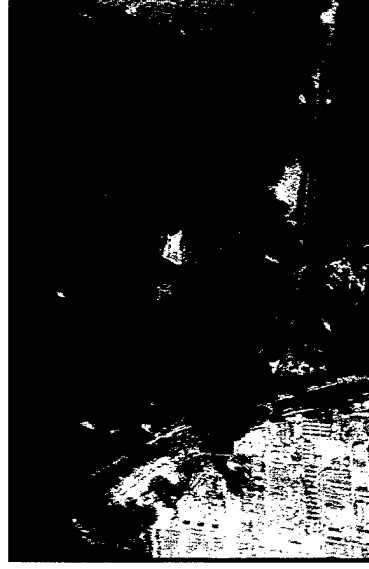
Original Image:
Airfield 1 Scene 2 Band 2



Independent Spectral Bands
pSNR 25.83



DCT Across Spectral Bands
pSNR 30.24



KLT Across Spectral Bands
pSNR 32.61

6.4 Multiwavelet Signal Processing

A very recent development in wavelet theory is the class of multiwavelets – wavelet decompositions based on matrix dilation equations. While these structures have been developed by mathematicians, significant gaps in application of multiwavelets presented themselves to us. We have addressed the computational issues of pre-filtering and symmetric extension for multiwavelet signal and image processing, and applied multiwavelet filters to image compression and denoising [49], [105], [106]. We summarize these results here.

Multiwavelets are also based on a multiresolution analysis, however one that is based on several scaling functions. A basis for the coarse approximation space V_0 is generated by translates of N scaling functions $\phi_1(t-k)$, $\phi_2(t-k)$, \dots , $\phi_N(t-k)$. The vector $\Phi(t) = [\phi_1(t), \dots, \phi_N(t)]^T$, will satisfy a *matrix* dilation equation (analogous to the scalar case)

$$\Phi(t) = \sum_k C[k] \Phi(2t - k). \quad (44)$$

The coefficients $C[k]$ are N by N matrices instead of scalars.

Associated with these scaling functions are N wavelets $w_1(t), \dots, w_N(t)$, satisfying the *matrix* wavelet equation

$$W(t) = \sum_k D[k] \Phi(2t - k). \quad (45)$$

Again, $W(t) = [w_1(t), \dots, w_N(t)]^T$ is a vector and the $D[k]$ are N by N matrices.

As in the scalar case, one can find the conditions of orthogonality and approximation for multiwavelets [103, 104, 37, 76].

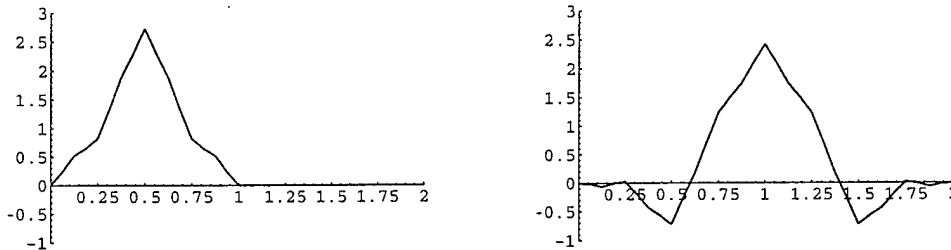


Figure 44: Geronimo-Hardin-Massopust pair of scaling functions.

A very important multiwavelet system was constructed by J. Geronimo, D. Hardin, and P. Massopust [32] (see [9] for another early multiwavelet construction). Their system contains the two scaling functions $\phi_1(t)$, $\phi_2(t)$ shown in Figure 44 and the two wavelets $w_1(t)$, $w_2(t)$ shown in Figure 45. The dilation and wavelet equations for this system have four coefficients:

$$\Phi(t) = \begin{bmatrix} \phi_1(t) \\ \phi_2(t) \end{bmatrix} = C[0]\Phi(2t) + C[1]\Phi(2t-1) + C[2]\Phi(2t-2) + C[3]\Phi(2t-3),$$

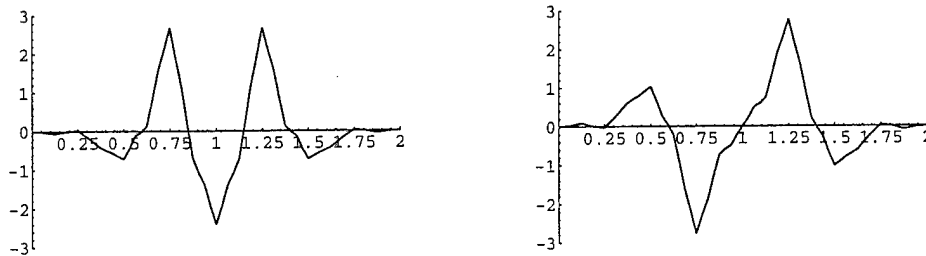


Figure 45: Geronimo-Hardin-Massopust multiwavelets.

$$W(t) = \begin{bmatrix} w_1(t) \\ w_2(t) \end{bmatrix} = D[0]\Phi(2t) + D[1]\Phi(2t - 1) + D[2]\Phi(2t - 2) + D[3]\Phi(2t - 3),$$

Actual values for the coefficient matrices $C[i]$ and $D[i]$ can be found in [105]. There are four remarkable properties of the Geronimo-Hardin-Massopust scaling functions:

- They each have short support (the intervals $[0, 1]$ and $[0, 2]$).
- Both scaling functions are symmetric, and the wavelets form a symmetric/anti-symmetric pair.
- All integer translates of the scaling functions are orthogonal.
- The system has second order of approximation (locally constant and locally linear functions are in V_0).

Let us stress that a scalar system with one scaling function cannot combine symmetry, orthogonality, and second order approximation. Moreover, a solution of a scalar dilation equation with four coefficients is supported on the interval $[0, 3]$!

Other useful constructions of multiwavelets are given in [9, 30, 107, 112, 77, 61].

6.4.1 Multiwavelets and multirate filter banks

Corresponding to each multiwavelet system is a matrix-valued multirate filter bank or multifilter. A multiwavelet filter bank [103] has “taps” that are $N \times N$ matrices (we work primarily with $N = 2$). The principal example is the 4-coefficient symmetric multiwavelet filter bank whose lowpass filter was reported in [32]. This filter is given by the four 2×2 matrices $C[k]$. The corresponding 2-channel, 2×2 matrix filter bank operates on *two* input data streams, filtering them into *four* output streams, each of which is downsampled by a factor of 2. This is shown in Figure 46. Each row

of the multifilter is a combination of two ordinary filters, one operating on the first data stream and the other operating on the second. For example, the first lowpass multiwavelet filter $C[0]$ operates as $c_{0,0}[k]$ on the first input stream and $c_{0,1}[k]$ on the second. It is a combination of the Haar filter $\{1, 1\}$ on the first stream and the unit impulse response on the second stream.

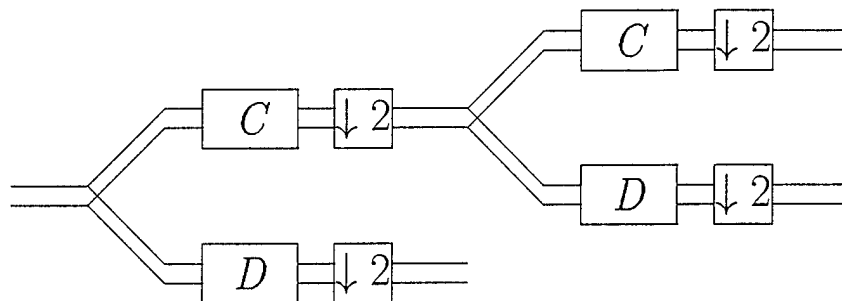


Figure 46: A multiwavelet filter bank, iterated once.

The matrix filter coefficients must satisfy the orthogonality (“block-paraunitarity”) condition

$$\sum_{k=0}^{N-1} C[k] C[k-2l]^T = 2\delta_{0,l} I. \quad (46)$$

The lowpass filter C and highpass filter D consist of coefficients corresponding to the dilation equation (44) and wavelet equation (45). But in the multiwavelet setting these coefficients are n by n matrices, and during the convolution step they must multiply vectors (instead of scalars). This means that multifilter banks *need n input rows*. We developed several ways to produce those rows. The first method, an oversampled scheme, consisted of simply repeating each row of data twice to produce the two input rows required. This method yielded good results at denoising one-dimensional signals, but performed poorly for compression, not a surprising outcome. Our second method, which was successful at both compression and denoising, exploited the approximation properties of the multiwavelet scaling functions to *prefilter* each row of input data, resulting in two new rows with half the number of data points for input to the multifilter. This later method had the advantage of preserving critical sampling, so that an input dataset with N points remained of size N after prefiltering. Another advantage of this approximation-based preprocessing method is that it fits naturally with symmetric extension for multiwavelets (discussed below). In other words, if we symmetrically extend a finite length signal $f[n]$ at its boundaries and implement the approximation formulas, then the two rows output by the preprocessor will have the appropriate symmetry.

In the setting of purely two-dimensional signal processing, we described an additional algorithm for multiwavelet filtering (two rows at a time), and developed a new

family of multiwavelets (the constrained pairs) that is well-suited to this two-row-at-a-time filtering. Further details appear in [105].

In practice all signals have finite length, so we must devise techniques for filtering such signals at their boundaries. There are two common methods for filtering at the boundary that preserve critical sampling. The first is circular periodization (periodic wrap) of the data. This method introduces discontinuities at the boundaries; however, it can be used with almost any filter bank. The second approach is symmetric extension of the data. Symmetric extension preserves signal continuity, but can be implemented only with linear-phase (symmetric and/or antisymmetric) filter banks [95, 12, 56, 15]. We have developed symmetric extension for linear-phase multiwavelet filters, such as the Geronimo-Hardin-Massopust multifilters. This has proven useful for image compression applications.

Recall the basic problem: given an input signal $f[n]$ with N samples and a linear-phase (symmetric or antisymmetric) filter, how can we symmetrically extend f before filtering and downsampling in a way that preserves the critically sampled nature of the system? The possibilities for such an extension have been enumerated in [15]. Depending on the parity of the input signal (even- or odd-length) and the parity and symmetry of the filter, there is a specific non-expansive symmetric extension of both the input signal and the subband outputs. For example, an even-length input signal passed through an even-length symmetric lowpass filter should be extended by repeating the first and last samples, i.e., a half-sample symmetric signal is matched to a half-sample-symmetric filter. Similarly, when the lowpass filter is of odd length (whole-sample-symmetry), the input signal should be extended without repeating the first or last samples.

Each row of the GHM multifilter is a linear combination of two filters, one for each input stream. One filter (applied to the first stream) is of even length; the second is of odd length. Thus we extend the first stream using half-sample-symmetry (repeating the first and last samples) and extend the second stream using whole-sample-symmetry (*not* repeating samples). Then, when synthesizing the input signal from the subband outputs, we must symmetrize the subband data differently depending on whether it is going into an even- or odd-length filter. This approach has been extended to multifilters lacking linear-phase symmetry as well. Details appear in [105]; the upshot is that we have obtained a non-expansive transform of finite-length input data which behaves well at the boundaries under lossy quantization.

6.4.2 Denoising by soft thresholding

We have compared the numerical performance of GHM and constrained multiwavelets with Daubechies D_4 scalar wavelets. D_4 wavelets were chosen because they have two vanishing moments, are orthogonal and have four coefficients in the dilation equation — exactly like the GHM and constrained pairs. We perform these comparisons in

two standard wavelet applications: signal denoising and data compression. First we discuss the denoising.

Suppose that a signal of interest f has been corrupted by noise, so that we observe a signal g :

$$g[n] = f[n] + \sigma z[n],$$

where $z[n]$ is unit-variance, zero-mean Gaussian white noise. What is a robust method for recovering f from the samples $g[n]$ as best as possible? Donoho and Johnstone [28, 29] have proposed a solution via wavelet shrinkage or soft thresholding in the wavelet domain. Wavelet shrinkage works as follows:

1. Apply the cascade algorithm to get the wavelet coefficients corresponding to $g[n]$.
2. Choose a threshold $t_n = \sqrt{2 \log(n)} \gamma \sigma / \sqrt{n}$ and apply (soft) thresholding to the wavelet coefficients.
3. Invert the cascade algorithm to get the denoised signal $\hat{f}[n]$.

Donoho and Johnstone's algorithm offers the advantages of smoothness and adaptation. Wavelet shrinkage is *smooth* in the sense that the denoised estimate \hat{f} has a very high probability of being as smooth as the original signal f , in a variety of smoothness spaces (Sobolev, Hölder, etc.). Wavelet shrinkage also achieves near-minimax mean-square-error among possible denoisings of f , measured over a wide range of smoothness classes. In these numerical senses, wavelet shrinkage is superior to other smoothing and denoising algorithms. Heuristically, wavelet shrinkage has the advantage of not adding "bumps" or false oscillations in the process of removing noise, because of the local and smoothness-preserving nature of the wavelet transform. Wavelet shrinkage has been successfully applied to SAR imagery as a method for clutter removal [74]. It is natural to attempt to use multiwavelets as the transform for a wavelet shrinkage approach to denoising, and compare the results with scalar wavelet shrinkage.

We implemented Donoho's wavelet shrinkage algorithm using several additional remarks from [74]. We compared the performance of the D_4 scalar wavelet transform with oversampled and critically sampled multiwavelet schemes. In the oversampled scheme, the first row is multiplied by $\sqrt{2}$, to better match the first eigenvector of the GHM system. The critically sampled scheme obtains two input rows $v_{1,n}$, $v_{2,n}$ from a single row of data. After reconstruction the two output rows $\hat{v}_{1,n}$, $\hat{v}_{2,n}$ are deapproximated to yield the output signal $\hat{f}[n]$. Boundaries are handled by symmetric data extension for the critically sampled (approximation/deapproximation) and oversampled schemes, and by circular periodization for D_4 .

Combining these various techniques, we were able to apply multiwavelet denoising to imagery. We added white Gaussian noise to the *Lenna* image, and applied three wavelet transforms for denoising by wavelet shrinkage: the GHM multiwavelet with approximation, GHM with repeated row, and the Daubechies 4-tap scalar wavelet. The experimental results are shown in Table 4 and the resulting images in Figure 47. The GHM multiwavelet with approximation was superior to D_4 both numerically and subjectively; the approximation-based preprocessing seemed to reduce the Cartesian artifacts present in the scalar wavelet shrinkage. This can be seen, for example, in the facial features (eyes, nose) of the *Lenna* images shown in Figure 47. The GHM-repeated row scheme suffered because we had to repeat rows in first the x dimension and then in the y dimension, altering the correlations of the data. This produces the broad stripes in the image denoised with the repeated row scheme.

Figure 47: See attached page

	Noise	GHM with approximation	GHM with repeated row	D_4
ℓ^1 error	19.93	6.87	9.70	8.09
ℓ^2 error	24.98	9.75	12.6	11.4

Table 4: Denoising of *Lenna* image via wavelet-shrinkage.

6.4.3 Transform-based image coding

One of the most successful applications of the wavelet transform is image compression. A transform-based coder operates by transforming the data to remove redundancy, then quantizing the transform coefficients (a lossy step), and finally entropy coding the quantizer output. Because of their energy compaction properties and correspondence with the human visual system, wavelet representations have produced superior objective and subjective results in image compression [64], [130], [11], [16]. Since a wavelet basis consists of functions with short support (for high frequencies) and long support (for low frequencies), large smooth areas of an image may be represented with very few bits, and detail added where it is needed. Multiwavelet decompositions offer all of these traditional advantages of wavelets, as well as the combination of orthogonality, short support, and symmetry. The short support of multiwavelet filters limits ringing artifacts due to subsequent quantization. Symmetry of the filter bank both leads to efficient boundary handling and preserves centers of mass, lessening the blurring of fine-scale features. Orthogonality is useful because it means that rate-distortion optimal quantization strategies may be employed in the transform domain



Lenna image with Gaussian noise
MSE 24.98



GHM-with-approximation
multiwavelet denoising, MSE 9.75



Daubechies 4 scalar
wavelet denoising, MSE 11.4



GHM-repeated-row
multiwavelet denoising, MSE 12.6

Figure 47: Denoising comparison.

and still lead to optimal time-domain quantization (at least when error is measured in a mean-square sense). Thus it is natural to consider the use of multiwavelets in a transform-based image coder.

We compared the new two-dimensional multiwavelet algorithms with a D_4 scalar wavelet in a production image coding system. Five types of wavelet transform were used:

- D_4 scalar wavelet
- Approximation/deapproximation preprocessing with GHM multiwavelets
- Adjacent rows input with GHM multiwavelets
- Adjacent rows input with symmetric pair
- Adjacent rows input with two different constrained pairs

Each of these wavelet transforms was followed by entropy-constrained scalar quantization and entropy coding. We made the assumption that the histograms of subband (or wavelet transform subblock) coefficient values obeyed a Laplacian distribution [64], and designed a uniform scalar quantizer. The quantizer optimized the bit allocation among the different subbands by using an operational rate-distortion approach (minimizing the functional $D + \lambda R$) [93]. We then entropy-coded the resulting coefficient streams using a combination of zero-run-length coding and adaptive Huffman coding, as in the FBI's Wavelet Scalar Quantization standard [7].

Compression Ratio	8:1	16:1	32:1	64:1
	pSNR	pSNR	pSNR	pSNR
Daubechies 4	35.6	32.3	29.3	26.8
GHM with appr./deappr.	35.3	31.8	29.4	27.1
Adjacent Row Processing:				
GHM	24.1	21.3	19.7	18.4
symmetric pair	31.1	27.3	24.0	21.8
constrained pair #1	32.4	28.5	25.1	23.0
constrained pair #2	31.9	28.2	25.0	22.8

Table 5: Peak SNRs for compression of *Lenna*.

We applied these different wavelet image coders to the *Lenna* (NITF6) image, as well as a geometric test pattern, at a variety of compression ratios. The results are shown in Tables 5 and 6, and in Figures 48 and 49. On *Lenna*, the GHM multiwavelet with approximation mildly outperformed the D_4 scalar wavelet at compression ratios of 32:1 and 64:1. The images in Figure 48 show that the GHM-approximation scheme

preserves more texture in the hat and, as in the denoising application, produce fewer Cartesian artifacts than the scalar wavelet scheme. The repeated-row method did not work well on *Lenna*. However, the repeated-row method produced the best compressions of the test pattern image (49) at intermediate compression ratios (16:1 and 32:1), with the constrained pair #1 “CP-1” outperforming both D_4 and GHM with approximation. When using the repeated row algorithm, the constrained pairs significantly outperformed the GHM symmetric multiwavelet, demonstrating the importance of the eigenvector constraints used in definition of the constrained pairs. A close look at the details of the compressed/decompressed test patterns shows that the CP-1 compression did a better job of preserving the checkerboard and “rang” over a shorter distance than the D_4 compression. The alteration of the checkerboard pattern in the D_4 compression may be due to the lack of linear-phase symmetry in the wavelet filters.

These preliminary results suggest that multiwavelets are worthy of further investigation as a technique for image compression. Issues to address include the design of multiwavelets with symmetry and higher order of approximation than the GHM system, the role of eigenvector constraints, and also further exploration of regularity for multiwavelets [119]. One might also apply zerotree-coding methods [92] in a multiwavelet context.

Compression Ratio	8:1	16:1	32:1	64:1
	pSNR	pSNR	pSNR	pSNR
Daubechies 4	48.5	31.4	23.0	19.8
GHM with appr./deappr.	52.4	34.0	18.3	16.8
Adjacent Row Processing:				
GHM	29.8	25.4	20.1	15.8
symmetric pair	33.3	28.3	20.9	16.5
constrained pair 1	42.2	32.3	23.9	19.0
constrained pair 2	33.3	30.2	21.6	17.6

Table 6: Peak SNRs for compression of geometric test pattern.

Figure 48: See attached page

Figure 49: See attached page



Lenna original



**GHM approximation-based
multiwavelet compression
64:1, pSNR 27.1**

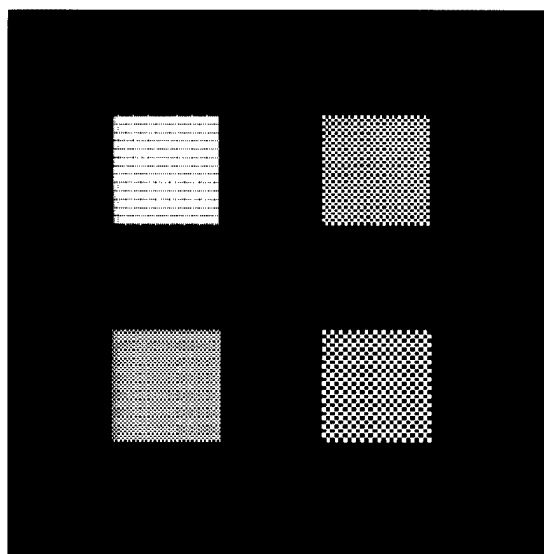


**Daubechies 4
scalar wavelet compression
64:1, pSNR 26.8**

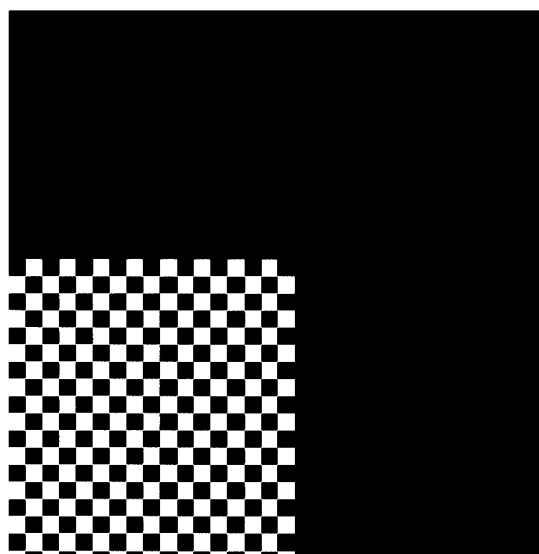


**Constrained-Pair #1
multiwavelet compression
64:1, pSNR 23.0**

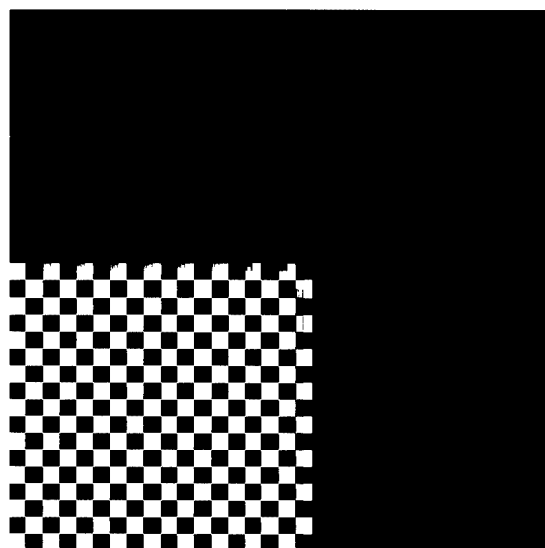
Figure 48: Lenna compression comparison.



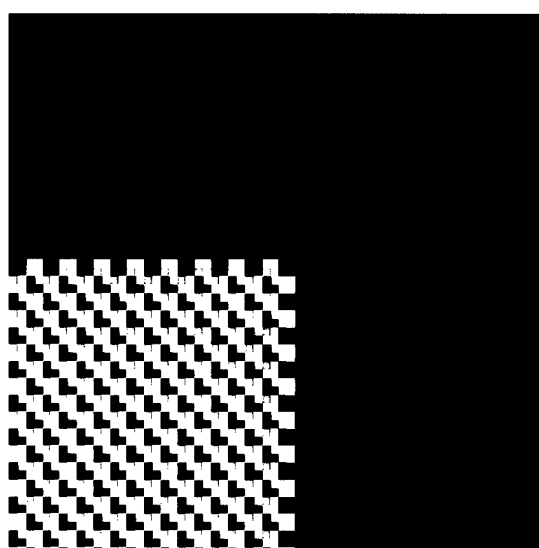
Original geometric pattern



**Detail of original pattern
(corner of lower left checkerboard)**



**Detail of CP-1 multiwavelet
compression (32:1, pSNR 23.9)**



**Detail of D4 scalar wavelet
compression (32:1, pSNR 23.0)**

Figure 49: Pattern compression comparison.

In summary, we have applied the new mathematical constructions of multiwavelets. Multiwavelets offer the advantages of combining symmetry, orthogonality, and short support, properties not mutually achievable with scalar 2-band wavelet systems. However, multiwavelets differ from scalar wavelet systems in requiring two or more input streams to the multiwavelet filter bank. We described two methods (repeated row and approximation/deapproximation) for obtaining such a vector input stream from a one-dimensional signal. We developed the theory of symmetric extension for multiwavelet filter banks, which matches nicely with approximation-based preprocessing. We then applied this arsenal of techniques to two basic signal processing problems, denoising via thresholding (wavelet shrinkage), and data compression. After developing the approach via model problems in one dimension, we applied the various new multiwavelet approaches to the processing of images, frequently obtaining performance superior to the comparable scalar wavelet transform. These results suggest that further work in the design and application of multiwavelets to signal and image processing is well warranted.

7 Technical Conclusions

In conclusion, this project explored and delineated the theory of a broad new class of mathematical transforms, the rank M wavelets. We developed a production-quality software package (WaveTool) for the design and implementation of these transforms, and applied them to concrete problems in signal processing and communications. In the theoretical realm, we discovered complete parametrizations of the family of rank M wavelets, and developed and applied mathematical tools for measuring the Sobolev smoothness of rank M wavelet systems, leading to surprising asymptotic regularity results. We also developed numerous techniques for rank M wavelet filter design, including those based on approximation (vanishing moments) and on smoothness of the iterated filter (regularity). Methods for the construction of full rank M wavelet matrices led directly to fast algorithms for computation, particularly in the case of the cosine-modulated filter banks.

This project also saw the design and completion of a broad and flexible software system, "WaveTool", for prototyping wavelet algorithms. After installation at a number of government, academic, and industrial beta sites, this software was successfully turned into a commercial product. Finally, new applications of rank M wavelets and design techniques were explored, particularly to multicarrier modulation for broadband communications. This latter application has led directly to a chipset product for high-bitrate communications across twisted-pair copper wire and hybrid fiber-coax networks. We also improved our leading-edge wavelet image compression algorithms, and applied wavelet-based compression to sonar, seismic, and multispectral image data. Finally, we produced the first significant applications of the new multiwavelet

techniques to signal and image processing, yielding promising results in denoising and compression.

Directions for further research that were suggested by this work include:

- The systematic application of wavelet-based smoothness measures in image quantization and compression. Initial steps have been taken in [19].
- Deeper exploration of cosine-modulated filter bank and wavelet structures, including filter design and fast algorithms for computation [71], [42].
- Thorough development of wavelet-based systems for multicarrier modulation, addressing issues such as equalization and system latency [47].
- Deeper exploration of multiwavelet techniques for signal processing, including advances in multiwavelet filter design and preprocessing algorithms [129].

Of course this is only a partial list; many other topics can be named.

8 Participants

Employees at Aware who participated in work on this contract included: Dr. Peter Niels Heller, Dr. Howard L. Resnikoff, Dr. Michael Tzannes, Dr. John Weiss, Dr. Edmund Reiter, Hemant Singh, Lev Weisfeiler, Vasily Strela, W. Knox Carey, Dr. Richard Gross, Dr. Stephen DelMarco, Karl Jagler, L. Scott Hills, and Anna Rounbehler.

A number of consultants, each of them experts in their respective fields, were employed as well. They included: Professor R. O. Wells, Jr. of Rice University (wavelet mathematics and computation), Professor P. P. Vaidyanathan of Cal Tech (multi-rate signal processing), Professor Truong Q. Nguyen of the University of Wisconsin (design of wavelet and filter bank structures), Dr. Ramesh Gopinath of Rice University (wavelets and signal processing), and Dr. Sundar Narasimhan of MIT (graphics software for WaveTool).

9 Publications

Publications arising from the work completed under this contract include five refereed journal articles, 14 conference papers, and two book chapters. Specifically, they include:

Journal Articles:

1. P. Steffen, P. N. Heller, R. A. Gopinath, C. S. Burrus, "Theory of Regular M -band Wavelets," *IEEE Trans. on SP*, 41 (1993), pp. 3497-3511.
2. P. N. Heller, "Rank M Wavelets With N Vanishing Moments," *SIAM J. Matrix Analysis*, 16 (1995), pp. 502-519.
3. H. L. Resnikoff, "Analytic Representation of Compactly Supported Wavelets," in *Festschrift for Hans Bremermann, BioSystems*, 34 (1995), pp. 259-272.
4. V. Strela, P. N. Heller, G. Strang, P. Topiwala, and C. Heil, "The Application of Multiwavelet Filter Banks to Image Processing," to appear, *IEEE Trans. on Image Processing*.
5. P. N. Heller and R. O. Wells, Jr., "Sobolev Regularity for Rank M Wavelets," submitted to *SIAM J. Math. Analysis*, 1996.

Conference Papers:

1. P. N. Heller and H. L. Resnikoff, "Regular M -band wavelets and applications," *Proc. IEEE ICASSP*, Minneapolis, 1993.
2. C. Bosman and E. C. Reiter, "Seismic data compression using wavelet transforms," *SEG Annual Meeting Extended Abstracts*, pp. 1261-1264, 1993.
3. P. N. Heller and K. Jagler, "Wavelet compression of multispectral imagery," in *Proc. Industry Workshop - Data Compression Conference*, Snowbird, Utah, 1994.
4. H. L. Resnikoff, "Perfect reconstruction and wavelet matrix windows for harmonic analysis," in *Proc. SPIE*, San Diego, CA, 1994.
5. P. N. Heller, "Lagrange M -th band filters and the construction of smooth M -band wavelets," in *Proc. IEEE-SP Intl. Symp. on Time-Frequency and Time-Scale Analysis*, Philadelphia, PA, 1994, pp. 108-111.
6. E. C. Reiter and P. N. Heller, "Wavelet transform based compression of NMO-corrected CDP gathers," *Society of Exploration Geophysicists 64th Annual Mtg.*, Los Angeles, 1994., pp. 731-734.
7. M. A. Tzannes, M. C. Tzannes, J. Proakis, P. N. Heller, "DMT Systems, DWMT systems, and digital filter banks," in *Proc. IEEE ICC*, New Orleans, LA, 1994.

8. P. N. Heller, J. M. Shapiro, and R. O. Wells, Jr., "Optimally smooth symmetric quadrature mirror filters for image coding," in *Proc. SPIE 2491, Wavelet applications for dual use*, Orlando, FL, April, 1995.
9. P. N. Heller, V. Strela, G. Strang, P. Topiwala, C. Heil, L. S. Hills, "Multi-wavelet filter banks for data compression," in *Proc. IEEE ISCAS '95*, Seattle, Washington.
10. P. N. Heller, T. Q. Nguyen, H. Singh, W. K. Carey, "Linear-Phase M -band wavelets with application to image coding," in *Proc. IEEE ICASSP*, Detroit, MI, 1995.
11. S. DelMarco, P. N. Heller, and J. Weiss, "An M -band, 2-dimensional translation-invariant wavelet transform and applications," in *Proc. IEEE ICASSP*, Detroit, MI, 1995.
12. H. Singh and P. N. Heller, "WaveTool: an integrated software for wavelet and multirate signal processing," *IEEE Intl. Conf. on Image Proc.*, Washington, D.C., 1995.
13. V. Strela, P. N. Heller, G. Strang, P. Topiwala and C. Heil, "Application of multiwavelets to signal and image processing," in UK Symposium on Applications of Time-Frequency and Time-Scale Methods, Warwick, UK, 1995.
14. M. Lang and P. N. Heller, "The Design of maximally smooth wavelets," *IEEE ICASSP*, Atlanta, GA, 1996.

Book Chapters:

1. P. N. Heller and R. O. Wells, Jr., "The Spectral Theory of Multiresolution Operators and Applications," in *Wavelets: Theory, Algorithms, and Applications*, C. K. Chui, L. Montefusco, L. Puccio, eds., Academic Press, San Diego, 1994, pp. 13-32.
2. P. N. Heller, "Tutorial 2.5: Subband and Wavelet Transforms: Theory, Design, and Applications: Educational Software," in *Microsystems Technology for Multimedia Applications*, IEEE ISCAS Tutorial volume, 1995.

In addition a number of invited talks were given, including:

- Nordic Postgraduate course on Wavelets and Filter Banks, Helsinki, Finland, 1994 – P. N. Heller
- Sherman Memorial Lecture at Indiana University, Bloomington, IN, 1994 – H. L. Resnikoff

- Presentation at the 1993 Taormina, Italy conference on Wavelets and Applications – R. O. Wells, Jr.
- Hour talk at Argonne Workshop on Wavelets and Large-Scale Image Processing, Argonne, IL, 1994 – P. N. Heller

10 Transition of Technology to Government and Commercial Uses

10.1 WaveTool

The WaveTool wavelet and multirate design and algorithm prototyping software was installed at the following government, industrial, and academic “beta sites”:

- National Security Agency (R5): Adolf Cusmariu and Mark Marson
- National Agency: Laszlo Fulop
- Analog Devices, Norwood, MA
- General Instrument, Hatboro, PA
- Massachusetts Institute of Technology (Math Dept. and Civil Engineering Dept.)
- California Institute of Technology (Electrical Engineering Dept.)
- Rice University (Electrical and Computer Engineering Dept. and Math Dept.)
- University of California at Davis (Electrical Engineering Dept.)
- Helsinki University of Technology (Electrical Engineering Dept.)

Feedback from these users proved very useful for later improvements to WaveTool, preceding its release as a commercial product in April 1995. Since then, a number of copies of this specialized signal processing software have been delivered to industry and academic customers around the world.

10.2 Other commercial products

Work done on this contract has contributed to numerous other Aware products in both last-mile telecommunications and in image compression. In telecommunications, the wavelet filter designs and computational algorithms of section 4.4 as applied to multicarrier modulation (section 6.1) have been designed into an application-specific

integrated circuit (ASIC) with our partner, Analog Devices. This chip, the AD6434, implements a 384-tone DWMT algorithm using a cosine-modulated wavelet filter bank with rank $M = 384$ and genus $g = 6$. It is intended for use in both hybrid fiber-coax and twisted-pair systems. As part of a joint project with DSC Communications Corp., the AD6434 has been incorporated into DSC's MediaSpan architecture to provide the cable telephony subsystem. Just recently, the first successful phone call over this DWMT-based telephony system was made. Similar DWMT algorithms are being designed into an Aware/ADI chipset for a VDSL system (high-bitrate communications over twisted-pair copper lines of length 1000 to 3000 feet) that will be released during 1997.

This project's advances in wavelet filter design and image compression algorithms have contributed to Aware's numerous successful products for wavelet-based still image compression. The extension of wavelet methods to seismic data compression has led directly to a software product, SeisPact, that is seeing wide use. SeisPact has been widely deployed for ship-to-shore communication of seismic exploration results via Inmarsat (satellite) links, as discussed in section 6.2.4. SeisPact is also being used in on-shore seismic data processing applications to minimize data storage requirements. Aware's AccuPress software for compression of 8-bit and 24-bit images derived a number of improvements from the project being reported, and is being used for a variety of applications from medical image compression to multimedia. In addition, Aware has become the leading vendor of compliance-certified implementations of the FBI's Wavelet Scalar Quantization [7] wavelet-based fingerprint compression algorithm.

10.3 Other government work

Some of the ideas explored during this contract have led to further government work as well. Aware has successfully investigated compression of one-dimensional acoustic data via a Phase I SBIR with the Navy [26], and is currently working on Phase II of the same project, intended to lead to a compression product for one-dimensional data.

We also successfully pursued and completed a Phase I SBIR for the National Institute of Standards and Technology, the "Wavelet Image Compression Workbench," [40] that took many of the ideas from this contract involving choice of wavelet basis and multiresolution tree and incorporated them into a software plug-in module for Adobe Photoshop. This led directly to a software workbench for comparing and prototyping wavelet image compression algorithms, and to a compression product that we are selling today.

Finally, new applications of wavelet techniques to communications were the subject of a follow-on government agency contract with Howard L. Resnikoff [85].

References

- [1] —, "Wavelet-Unified Design Toolbox for Digital Signal Processing: Technical Report #1," Aware Technical Report AD930309, 1993.
- [2] —, "Wavelet-Unified Design Toolbox for Digital Signal Processing: Annual Technical Report for 1993," Aware Technical Report DN04-07-03.R0, 1994.
- [3] —, "Wavetool Reference Manual," Aware Technical Document AD940228, 1994.
- [4] —, "Wavetool Users' Guide," Aware Technical Document AD940315, 1994.
- [5] "Coding of Moving Pictures and Associated Audio for Digital Storage Media up to about 1.5 MBit/s", *International norm MPEG1*, ISO No. 11172-3, 1992
- [6] "Asymmetric Digital Subscriber Line Transmission Standard," ANSI document T1E1.4/95-007R2.
- [7] Fed. Bureau of Investig., *WSQ Gray-Scale Fingerprint Image Compression Specification*, IAFIS-IC-0110-v2, Feb. 1993. Drafted by T. Hopper, C. Brislawn, and J. Bradley.
- [8] E. Adelson, E. Simoncelli, and R. Hingorani, "Orthogonal pyramid transforms for image coding," in *Proc. SPIE Visual Commun. and Image Proc. II*, Cambridge, MA 1987.
- [9] B. Alpert, "A class of bases in L^2 for the sparse representation of integral operators," *SIAM J. Math. Analysis*, vol. 24, 1993.
- [10] R. Ansari, C. Guillemot, and J. F. Kaiser, "Wavelet construction using Lagrange halfband filters," *IEEE Trans. on Circuits and Systems*, 38 (1991), pp. 1116-1118.
- [11] M. Antonini, M. Barlaud, P. Mathieu, and I. Daubechies, "Image coding using the wavelet transform," *IEEE Trans. on Image Processing*, vol. 1, pp. 205-220, 1992.
- [12] R. H. Bamberger, S. L. Eddins, and V. Nuri, "Generalized symmetric extension for size-limited multirate filter banks," *IEEE Trans. on Image Proc.*, vol. 3, pp. 82-86, 1994.
- [13] R. Blahut, *Theory of Error-Correcting Codes*, Addison-Wesley, Reading, MA, 1994.
- [14] C. Bosman and E. Reiter, "Seismic data compression using wavelet transforms," *SEG Annual Meeting Extended Abstracts*, pp. 1261-1264, 1993.
- [15] C. Brislawn, "Classification of symmetric wavelet transforms," Los Alamos Technical Report, 1993.
- [16] B. V. Brower, "Low-bit-rate image compression evaluations," *Proc. SPIE*, Orlando, FL, April 4-9, 1994.

- [17] W. K. Carey and P. N. Heller, "Dead zone quantizers in wavelet image compression," Aware Technical report XAD940822, Cambridge, MA, 1994.
- [18] W. K. Carey and P. N. Heller, "Uniform quantization and the generalized Gaussian," Aware Technical report XAD940815, Cambridge, MA, 1994.
- [19] W. K. Carey, S. Hemami, and P. N. Heller, "Smoothness-constrained wavelet image compression," *Proc. IEEE Intl. Conf. on Image Proc.*, Lausanne, 1996.
- [20] A. Cohen and J. P. Conze, "Régularité des bases d'ondelettes et mesures ergodiques," *Rev. Math. Iberoamericana*, vol. 8 (1992), pp. 351-365.
- [21] A. Cohen and I. Daubechies, "Non-separable bidimensional wavelet bases," *Rev. Math. Iberoamericana*, vol. 9 (1993).
- [22] R. Coifman and M. V. Wickerhauser, "Entropy-based algorithms for best basis selection," *IEEE Trans. on Info. Theory*, vol 38 (1992), pp. 713-718.
- [23] I. Daubechies, "Orthonormal bases of compactly supported wavelets", *Comm. Pure Appl. Math.*, 41 (1988), pp. 909-996.
- [24] I. Daubechies, *Ten Lectures on Wavelets*, CBMS Conference Series 61, SIAM, Philadelphia, 1992.
- [25] S. DelMarco et. al., "Real-time compression of sonar data for ASW simulation," Final Technical Report to ARPA, 1993 (also Aware technical report AD 931332).
- [26] S. DelMarco, Final Report for U. S. Navy Phase I SBIR contract N00421-95-C-1147, "M-band Acoustic Data Compression," Aware, Inc., Bedford, MA, 1996.
- [27] S. DelMarco, J. Weiss, and K. Jagler, "Sonar transient detection using the TI-wavelet transform," *Optical Engineering*, to appear, 1994.
- [28] D. Donoho, "De-noising by soft-thresholding," *IEEE Trans. Inf. Theory*, vol. 41, pp. 613-627, 1995.
- [29] D. L. Donoho and I. M. Johnstone, "Ideal spatial adaptation via wavelet shrinkage," *Biometrika*, vol. 81, pp. 425-455, 1994.
- [30] G. Donovan, J. Geronimo, and D. Hardin, "Intertwining multiresolution analyses and the construction of piecewise polynomial wavelets," preprint, 1994.
- [31] T. Eirola, "Sobolev characterization of solutions of dilation equations," *SIAM J. Math. Analysis*, vol. 23 (1992), pp.1015-1030.

- [32] J. Geronimo, D. Hardin, and P. R. Massopust, "Fractal functions and wavelet expansions based on several functions," *J. Approx. Theory*, vol. 78, pp. 373-401, 1994.
- [33] H. Gharavi and A. Tabatabai, "Subband coding of monochrome and color images," *IEEE Transactions on Circuits and Systems*, vol. 35 (1988), pp. 207-214.
- [34] P. Gill, W. Murray, and M. Wright, *Practical Optimization*, Academic Press, San Diego, 1981.
- [35] R. A. Gopinath and C. S. Burrus, *Wavelet transforms and filter banks*, in *Wavelets: A Tutorial in Theory and Applications*, C. K. Chui, ed., Academic Press, San Diego, 1992, pp.603-654.
- [36] R. A. Gopinath and C. S. Burrus, "Theory of cosine-modulated wavelet bases and tight frames," *Appl. and Comp. Harmonic Analysis*, 1995.
- [37] C. Heil, G. Strang, and V. Strela, "Approximation by translates of refinable functions," to appear in *Numerische Mathematik*.
- [38] P. N. Heller, "Lagrange M -th band filters and the construction of smooth M -band wavelets," *Proc. IEEE-SP Intl. Symp. on Time-Frequency and Time-Scale Analysis*, Philadelphia, PA, 1994.
- [39] P. N. Heller, "Rank M wavelet matrices with N vanishing moments," *SIAM J. Matrix Analysis*, vol. 16 (1995), pp. 502-518.
- [40] P. N. Heller, Final Report for National Institute of Standards and Technology Phase I SBIR contract 50-DKNB-5-00156, "Wavelet Image Compression Workbench," Aware, Inc., Bedford, MA, 1995.
- [41] P. N. Heller and K. Jagler, "Wavelet compression of multispectral imagery," *Proc. Industry Workshop - Data Compression Conference*, Snowbird, Utah, 1994.
- [42] P. N. Heller, T. Karp, and T. Q. Nguyen, "A General formulation of modulated filter banks," submitted to *IEEE Trans. on SP*, 1996.
- [43] P. N. Heller, M. Lang, and R. O. Wells, Jr., "Smoothness criteria for wavelet filter design," in preparation for *IEEE Trans. on Signal Processing*, 1996.
- [44] P. N. Heller, T. Q. Nguyen, H. Singh, and W. K. Carey, "Linear-phase M -band wavelets with application to image coding," *Proc. IEEE ICASSP*, Detroit, MI, 1995.
- [45] P. N. Heller and H. L. Resnikoff, "Polynomials are generalized eigenfunctions of the wavelet transform," Aware Technical Report AD910912, 1991.

- [46] P. N. Heller, H. L. Resnikoff, and R. O. Wells, Jr., *Wavelet matrices and the representation of discrete functions*, in *Wavelets: A Tutorial in Theory and Applications*, C. K. Chui, ed., Academic Press, San Diego, 1992, pp. 15-50.
- [47] P. N. Heller, S. Sandberg, and M. A. Tzannes, "DWMT Equalization strategies for upstream VDSL," ANSI T1E1.4 VDSL Study Project document no. 96-205, July, 1996.
- [48] P. N. Heller, J. M. Shapiro, and R. O. Wells, Jr., "Optimally smooth symmetric quadrature mirror filters for image coding", in *Proc. SPIE 2491, Wavelet applications for dual use*, Orlando, FL, April, 1995.
- [49] P. Heller, V. Strela, G. Strang, P. Topiwala, C. Heil, and L. Hills, "Multiwavelet filter banks for data compression," *Proc. IEEE ISCAS*, Seattle, WA, May 1995.
- [50] P. N. Heller and R. Tolimieri, "A general construction of $M \times 2M$ perfect reconstruction filter banks", Aware Technical Report AD921202, Cambridge, MA, 1992.
- [51] P. N. Heller and R. O. Wells, Jr., *Spectral theory of multiresolution operators and applications*, in *Wavelets: Theory, Algorithms, and Applications*, C. K. Chui, L. Montefusco, and L. Puccio, eds., Academic Press, San Diego, 1994, pp. 13-32.
- [52] P. N. Heller and R. O. Wells, Jr., "Sobolev regularity for rank M wavelets," submitted to *SIAM J. Math. Analysis*, 1996.
- [53] O. Herrmann, "On the approximation problem in nonrecursive digital filter design," *IEEE Trans. Circuit Theory*, 18 (1971), pp. 411-413.
- [54] B. Hirosaki, "An orthogonally multiplexed QAM system using the discrete Fourier transform," *IEEE Trans. on Communications*, 29 (1981), pp. 982-989.
- [55] N. Jayant and P. Noll, *Digital Coding of Waveforms*, Prentice-Hall, Englewood Cliffs, NJ, 1984.
- [56] H. Kiya, K. Nishikawa, and M. Iwahashi, "A Development of symmetric extension method for subband image coding," *IEEE Trans. on Image Proc.*, vol. 3, pp. 78-81, 1994.
- [57] R. D. Koilpillai and P. P. Vaidyanathan, "Cosine modulated FIR filter banks satisfying perfect reconstruction," *IEEE Trans. on SP*, pp. 770-783, April 1992.
- [58] M. Lang and P. N. Heller, "The Design of maximally smooth wavelets," in *Proc. IEEE ICASSP '96*, Atlanta, GA, 1996.

- [59] M. Lang and P. N. Heller, "Maximally smooth wavelet filter banks and their design", article in preparation for *IEEE Trans. on Signal Processing*, 1996.
- [60] W. Lawton, "Tight frames of compactly supported affine wavelets," *J. Math. Phys.*, vol. 31 (1990), pp.1898-1901.
- [61] W. Lawton, S. L. Lee, and Z. Shen, "An algorithm for matrix extension and wavelet construction," *Mathematics of Computation*, to appear.
- [62] Y.-P. Lin and P. P. Vaidyanathan, "Linear-phase cosine modulated maximally decimated filter banks with perfect reconstruction," *IEEE Trans. on SP*, vol. 43, pp. 2525-2539, November 1995.
- [63] B. Macq and J. Y. Mertes, "Optimization of linear multiresolution transforms for scene adaptive coding," *IEEE Trans. on Signal Proc.*, vol. 41, pp. 3568-3571, 1993.
- [64] S. Mallat, "A Theory for multiresolution signal decomposition: the wavelet representation," *IEEE Trans. PAMI*, vol. 11, pp. 674-693, 1989.
- [65] H. S. Malvar, "Extended lapped transforms: fast algorithms and applications", *IEEE Trans. on SP*, November 1992.
- [66] H. S. Malvar, *Signal Processing with Lapped Transforms*, Norwood, MA, Artech House, 1992.
- [67] F. Mintzer, "On half-band, third-band, and N -th band FIR filters and their design," *IEEE Trans. on ASSP*, 30 (1982), pp.734-738.
- [68] T. Q. Nguyen, "A Class of generalized cosine modulated filter bank", *Proc. IEEE ISCAS 92*, pp. 943-946.
- [69] T. Q. Nguyen, "Near-perfect-reconstruction pseudo-QMF banks", *IEEE Trans. on SP*, vol. 42, pp. 65-76, January 1994.
- [70] T. Q. Nguyen, "Digital filter bank design quadratic constrained formulation," *IEEE Trans. on SP*, vol. 43, pp. 2103-2108, September 95.
- [71] T. Q. Nguyen and P. N. Heller, "Biorthogonal cosine-modulated filter banks," *Proc. IEEE ICASSP*, Atlanta, GA, 1996.
- [72] T. Q. Nguyen and R. D. Koilpillai, "The theory and design of arbitrary-length cosine-modulated filter banks and wavelets, satisfying perfect reconstruction," *IEEE Trans. on SP*, vol. 44, pp. 473-483, March 1996.

- [73] T. Q. Nguyen and P. P. Vaidyanathan, "Two-Channel perfect-reconstruction FIR QMF structures which yield linear-phase analysis and synthesis filters," *IEEE Trans. on ASSP*, 37 (1989), pp. 676-690.
- [74] J. Odegard, M. Lang, H. Guo, R. Gopinath, and C. Burrus, "Nonlinear wavelet processing for enhancement of images," *IEEE SP Letters*, 1995.
- [75] W. B. Pennebaker and J. L. Mitchell, *JPEG: Still Image Compression Standard*, New York, NY, Van Nostrand Reinhold, 1993.
- [76] G. Plonka, "Approximation order of shift-invariant subspaces of $L^2(R)$ generated by refinable function vectors", preprint, 1995.
- [77] G. Plonka, "Generalized spline wavelets", *Constr. Approx.*, to appear.
- [78] J. Proakis, *Digital Communication, 3rd edition*, McGraw Hill, New York, NY, 1994.
- [79] R. de Queiroz, T. Nguyen, and K. Rao, "Generalized linear-phase lapped orthogonal transforms," *Proc. IEEE ISCAS*, London, 1994.
- [80] K. Ramchandran and M. Vetterli. "Best wavelet packets in a rate-distortion sense," *IEEE Transactions on Image Processing*, vol. 2 (1993), pp. 160-175.
- [81] T. A. Ramstad, J. P. Tanem, "Cosine modulated analysis synthesis filter bank with critical sampling and perfect reconstruction", *Proc. IEEE ICASSP 91*, pp. 1789-1792, Toronto, 1991.
- [82] K. R. Rao and P. Yip, *Discrete Cosine Transform: Algorithms, Advantages, Applications*, Academic Press, New York, 1990.
- [83] E. C. Reiter and P. N. Heller, "Wavelet transform based compression of NMO-corrected CDP gathers," Society of Exploration Geophysicists 64th Annual Mtg., Los Angeles, 1994., pp. 731-734.
- [84] H. L. Resnikoff, "Perfect reconstruction and wavelet matrix windows for harmonic analysis," in *Proc. SPIE*, San Diego, CA, 1994.
- [85] H. L. Resnikoff, Final Report for government contract, 1996.
- [86] O. Rioul, "Simple regularity criteria for subdivision schemes," *SIAM J. Math Analysis*, vol. 23 (1992), pp.1544-1576.
- [87] O. Rioul and M. Vetterli, "Wavelets and signal processing", *IEEE Signal Processing Magazine*, October, 1991.

- [88] A. D. Rizos, J. G. Proakis, and T. Q. Nguyen, "Comparison of DFT and cosine-modulated filter banks in multicarrier modulation," *Proc. IEEE Intl. Conf. on Communications*, New Orleans, LA, 1994.
- [89] M. Russell, M. A. Tzannes, S. Sandberg, P. N. Heller, "CAP/DWMT Solution for VDSL with Passive/Active NT," ANSI T1E1.4 VDSL Study Project document no. 96-123, April, 1996.
- [90] S. D. Sandberg and M. A. Tzannes, "Overlapped discrete multitone modulation for high-speed copper wire Communications," *IEEE J. Select Areas on Communications*, 13 (1995), pp. 1571-1585.
- [91] K. Schittkowski, "NLPQL: A FORTRAN subroutine solving constrained nonlinear programming problems," *Ann. Oper. Res.*, vol. 5, pp. 485-500, 1986.
- [92] J. M. Shapiro, "Embedded image coding using zerotrees of wavelet coefficients," *IEEE Trans. on SP*, vol. 41, pp. 3445-3662, 1993.
- [93] Y. Shoham and A. Gersho, "Efficient bit allocation for an arbitrary set of quantizers," *IEEE Trans. ASSP*, vol. 36, pp. 1445-1453, 1988.
- [94] M. J. T Smith and T. P. Barnwell III, "Exact reconstruction techniques for tree-structured subband coders," *IEEE Trans. on ASSP*, 34 (1986), pp. 434-441.
- [95] M. J. T. Smith and S. Eddins, "Analysis-synthesis techniques for subband image coding," *IEEE Trans. ASSP*, vol. 38, pp. 1446-1456, 1990.
- [96] A. K. Soman and P. P. Vaidyanathan, "On orthonormal wavelets and paraunitary filter banks", *IEEE Trans. on Signal Processing*, 41 (1993), pp. 1170-1183.
- [97] A. Soman, P. P. Vaidyanathan, and T. Q. Nguyen, "Linear phase paraunitary filter banks," *IEEE Trans. on SP*, 41 (1993), pp. 3480-3496.
- [98] J. Stautner, "High quality audio compression for broadcast and computer applications," presented at SMPTE Advanced Television and Electronic Imaging Conference, San Francisco, Feb. 1992 (also Aware Technical Report AD920207).
- [99] J. Stautner, "Scalable audio compression for mixed computing environments," *93rd Conv. of the Audio Engineering Soc.*, San Francisco, CA, 1992.
- [100] P. Steffen, P. N. Heller, R. A. Gopinath, C. S. Burrus, "Theory of regular M -band wavelets," *IEEE Trans. on Signal Processing*, vol. 41 (1993), pp. 3497-3511.
- [101] G. Strang, "Wavelets and dilation equations," *SIAM Review*, 31 (1989), pp. 614-627.

- [102] G. Strang and T. Nguyen, *Wavelets and Filter Banks*, Wellesley, MA: Wellesley-Cambridge Press, 1995.
- [103] G. Strang and V. Strela, "Short wavelets and matrix dilation equations," *IEEE Trans. on SP*, vol. 43, pp. 108-115, 1995.
- [104] G. Strang and V. Strela, "Orthogonal multiwavelets with vanishing moments," *J. Optical Eng.*, vol. 33, pp. 2104-2107, 1994.
- [105] V. Strela, P. N. Heller, G. Strang, P. Topiwala, and C. Heil, "The Application of multiwavelet filter banks to image processing," to appear, *IEEE Trans. on Image Processing*.
- [106] V. Strela, P. N. Heller, G. Strang, P. Topiwala and C. Heil, "Application of multiwavelets to signal and image processing," in UK Symposium on Applications of Time-Frequency and Time-Scale Methods, Warwick, UK, 1995.
- [107] V. Strela and G. Strang, "Finite element multiwavelets," *Proc. NATO Conference, Maratea*, Boston, MA: Kluwer, 1995.
- [108] N. Tanabe and N. Farvardin, "Subband image coding using entropy-coded quantization over noisy channels," *IEEE Journal on Selected Areas in Communications*, vol. 10 (1992), pp. 926-942.
- [109] M. A. Tzannes and R. Gross, "An Example DWT waveform and measured performance of a DWT modem," ANSI T1E1.4 VDSL Study Project document no. 96-124, April, 1996.
- [110] M. A. Tzannes, M. C. Tzannes, J. Proakis, P. N. Heller, "DMT Systems, DWT systems, and digital filter banks," in *Proc. IEEE Intl. Conf. on Comm.*, New Orleans, LA, 1994.
- [111] M. A. Tzannes, M. C. Tzannes, and H. L. Resnikoff, "Discrete Wavelet Multitone Modulation," ANSI T1E1.4 contribution, 1993.
- [112] G. Uytterhoeven, *Multiwavelets for image compression*, Masters thesis, Katholieke Universiteit Leuven, May, 1994.
- [113] P. P. Vaidyanathan, "Quadrature mirror filter banks, M -band extensions and perfect-reconstruction techniques", *IEEE ASSP Magazine*, 4 (1987), pp. 4-20.
- [114] P. P. Vaidyanathan and T. Q. Nguyen, "Eigenfilters: A new approach to least squares FIR filter design and applications including Nyquist filters," *IEEE Trans. on Circuits and Systems*, pp. 11-23, January 1987.

- [115] P. P. Vaidyanathan and T. Q. Nguyen, "A 'TRICK' for the design of FIR half-band filters," *IEEE Trans. on Circuits and Systems*, 34 (1987), pp. 297-300.
- [116] P. P. Vaidyanathan, *Multirate Systems and Filter Banks*, Prentice Hall, Englewood Cliffs, NJ, 1992.
- [117] M. Vetterli and C. Herley, "Wavelets and filter banks," *IEEE Trans. on SP*, 40 (1992), pp. 2207-2233.
- [118] M. Vetterli and D. LeGall, "Perfect reconstruction FIR filter banks," *IEEE Trans. ASSP*, 37 (1989), pp. 1057-1071.
- [119] M. Vetterli and G. Strang, "Time-varying filter banks and multiwavelets," Sixth IEEE DSP Workshop, Yosemite, 1994.
- [120] L. Villemoes, "Energy moments in time and frequency for two-scale difference equation solutions and wavelets," *SIAM J. Math. Analysis*, vol. 23 (1992), pp. 1519-1543.
- [121] L. Villemoes, "Wavelet analysis of refinement equations," *SIAM J. Math. Analysis*, vol. 25 (1994), pp. 1433-1460.
- [122] H. Volkmer, "On the regularity of wavelets," *IEEE Trans. on Information Theory*, vol. 38 (1992), pp. 872-876.
- [123] H. Volkmer, "Asymptotic regularity of compactly supported wavelets," *SIAM J. Math. Analysis*, 1995.
- [124] S. B. Weinstein and P. M. Ebert, "Data transmission by frequency-division multiplexing using the discrete Fourier transform," *IEEE Trans. Comm. Technology*, 15 (1971), pp. 628-634.
- [125] G. Welland and M. Lundberg, "Construction of compact p -wavelets," *Constructive Approximation*, vol. 9 (1993), pp. 347-370.
- [126] R. O. Wells, Jr., "Parametrizing smooth compactly supported wavelets," *Trans. Amer. Math. Soc.*, 338 (1993), pp. 919-931.
- [127] G. Wornell, "Emerging applications of multirate signal processing and wavelets in digital communications," *Proceedings of the IEEE*, 84 (1996), pp. 586-603.
- [128] X.-G. Xia, J. Geronimo, D. Hardin, and B. Suter, "Computations of multi-wavelet transforms," *Proc. SPIE 2569*, San Diego, CA, July 1995.
- [129] X.-G. Xia, J. S. Geronimo, D. P. Hardin, and B. W. Suter, "Design of prefilters for discrete multiwavelet transforms," *IEEE Trans. on SP*, 44 (1996), pp. 25-35.

- [130] W. Zettler, J. Huffman, and D. Linden. "The application of compactly supported wavelets to image compression," *Proc. SPIE* **1244**, pp. 150-160, 1990.
- [131] H. Zou and A. H. Tewfik, "Discrete orthogonal M-band wavelet decompositions", *Proc. IEEE ICASSP-San Francisco*, March 1992.
- [132] E. Zwicker and H. Fastl, *Psychoacoustics*, Springer Verlag, New York, 1990.

All information contained herein is unclassified and released,
 except where noted otherwise.
 DATE 10-12-12 BY 100-12
 (AFSC)